MEMOIRS OF THE EUROPEAN MATHEMATICAL SOCIETY

Paul Biran
Octav Cornea
Jun Zhang

Triangulation, Persistence, and Fukaya Categories

MEMS Vol. 28 / 2026

EMS PRESS


EMS
PRESS

Memoirs of the European Mathematical Society

Edited by

Anton Alekseev (Université de Genève)
Hélène Esnault (Freie Universität Berlin)
Gerard van der Geer (Universiteit van Amsterdam)
Ari Laptev (Imperial College London)
Laure Saint-Raymond (Institut des Hautes Études Scientifiques)
Susanna Terracini (Università degli Studi di Torino)

The *Memoirs of the European Mathematical Society* publish outstanding research contributions in individual volumes, in all areas of mathematics and with a particular focus on works that are longer and more comprehensive than usual research articles.

The collection's editorial board consists of the editors-in-chief of the *Journal of the European Mathematical Society* and the *EMS Surveys in Mathematical Sciences*, along with editors of book series of the publishing house of the EMS as well as other distinguished mathematicians.

All submitted works go through a highly selective peer-review process.

Latest titles in this series:

M. Duerinckx, A. Gloria, *On Einstein's Effective Viscosity Formula*
R. Willett, G. Yu, *The Universal Coefficient Theorem for C*-Algebras with Finite Complexity*
B. Janssens, K.-H. Neeb, *Positive Energy Representations of Gauge Groups I. Localization*
S. Dipierro, G. Giacomin, E. Valdinoci, *The Lévy Flight Foraging Hypothesis in Bounded Regions*
A. Naor, *Extension, Separation and Isomorphic Reverse Isoperimetry*
N. Lerner, *Integrating the Wigner Distribution on Subsets of the Phase Space, a Survey*
B. Adcock, S. Brugiapaglia, N. Dexter, S. Moraga, *On Efficient Algorithms for Computing Near-Best Polynomial Approximations to High-Dimensional, Hilbert-Valued Functions from Limited Samples*
J. J. Carmona, K. Fedorovskiy, *Carathéodory Sets in the Plane*
D. Abramovich, Q. Chen, M. Gross, B. Siebert, *Punctured Logarithmic Maps*
T. Oh, M. Okamoto, L. Tolomeo, *Stochastic Quantization of the* Φ^3_3*-Model*
H. Sasahira, M. Stoffregen, *Seiberg–Witten Floer Spectra for* $b_1 > 0$
J. Kaad, D. Kyed, *The Quantum Metric Structure of Quantum* SU(2)
B. Zavyalov, *Almost Coherent Modules and Almost Coherent Sheaves*
J-P. Labesse, B. Lemaire, *La formule des traces tordue pour un corps global de caractéristique* $p > 0$
S. Ma, *Vector-Valued Orthogonal Modular Forms*
Y. Almog, B. Helffer, *On the Stability of Symmetric Flows in a Two-Dimensional Channel*
A. Bors, D. Panario, Q. Wang, *Functional Graphs of Generalized Cyclotomic Mappings of Finite Fields*
P. Schneider, O. Venjakob, *Reciprocity Laws for* (φ_L,Γ_L)*-Modules over Lubin–Tate Extensions*
A.-L. Dalibard, F. Marbach, J. Rax, *Linear and Nonlinear Parabolic Forward-Backward Problems*
L. Ertzbischoff, D. Han-Kwan, *On Well-Posedness for Thick Spray Equations*
J. Bettinelli, G. Miermont, *Compact Brownian Surfaces II. Orientable Surfaces*

Paul Biran
Octav Cornea
Jun Zhang

Triangulation, Persistence, and Fukaya Categories

EMS PRESS

Authors

Paul Biran
Department of Mathematics
ETH Zürich
Rämistrasse 101
8092 Zürich, Switzerland

Email: biran@math.ethz.ch

Octav Cornea
Département de Mathématiques et Statistique
Université de Montréal
Pavillon André-Aisenstadt
2920 Chemin de la tour, bureau 6151
Montréal QC H3T 1J8, Canada

Email: cornea@dms.umontreal.ca

Jun Zhang
Institute of Geometry and Physics
University of Science and Technology of China
96 Jinzhai Road
230026 Hefei, P. R. China

Email: jzhang4518@ustc.edu.cn

Each volume of the *Memoirs of the European Mathematical Society* is available individually or as part of an annual subscription. It may be ordered from your bookseller, subscription agency, or directly from the publisher via subscriptions@ems.press.

ISSN 2747-9080, eISSN 2747-9099
ISBN 978-3-98547-107-2, eISBN 978-3-98547-607-7, DOI 10.4171/MEMS/28

The content of this volume is licensed under the CC BY 4.0 license, with the exception of the logo and branding of the European Mathematical Society and EMS Press, and where otherwise noted.

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the Internet at http://dnb.dnb.de.

Published by EMS Press, an imprint of the

European Mathematical Society – EMS – Publishing House GmbH
Institut für Mathematik
Technische Universität Berlin
Straße des 17. Juni 136
10623 Berlin, Germany

Email: info@ems.press

https://ems.press

© 2026 European Mathematical Society

Typesetting: Fernando Gómez, Bahía Blanca, Argentina
Printed in Germany
Printed on acid free paper

Abstract

This memoir introduces a new algebraic notion: that of a *triangulated persistence category* (TPC), which refines the notion of a triangulated category in the same sense that a persistence module refines that of a vector space. The spaces of morphisms of such a TPC are persistence modules, and the category is endowed with a class of weighted distinguished triangles. Under favourable conditions, we show that the derived Fukaya category admits a TPC refinement, and we apply this to deduce a global rigidity result for spaces of compact, exact Lagrangians in certain Liouville manifolds: we construct a metric on this space with intrinsic symplectic properties.

Keywords. triangulated category, persistence module, symplectic manifold, Lagrangian submanifold, Floer homology, Fukaya category

Mathematics Subject Classification (2020): 53D12 (primary); 53D37, 55N31 (secondary)

Acknowledgments. The third author is grateful to Mike Usher for useful discussions. We thank Leonid Polterovich for mentioning to us the work of Fan–Filip [30]. We thank the referee of the first version of this memoir for pointing out the relation between our algebraic construction and Verdier localization. We are also grateful to a second referee for a very careful and critical reading of a later version, for pointing out several inaccuracies, and for making very useful suggestions that helped to improve the exposition. We thank a third referee for pointing out the paper [51] (see also Remark 2.20).

Funding. The second author was supported by an individual NSERC Discovery grant. Part of this work was completed while the third author held a CRM–ISM Postdoctoral Research Fellowship at the *Centre de recherches mathématiques* in Montréal. He thanks the institute for its warm hospitality. The third author is partially supported by the National Key R&D Program of China (No. 2023YFA1010500), the NSFC (Nos. 12301081 and 12361141812), and the USTC Research Funds of the Double First-Class Initiative.

Contents

1 Introduction 1
- 1.1 Persistence and triangulation 3
- 1.2 TPC refinements of the Fukaya derived category 5

2 Triangulation persistence categories: Algebra 101 7
- 2.1 Triangular weights 7
- 2.2 Persistence categories 11
- 2.3 Triangulated persistence categories 21
- 2.4 A persistence triangular weight on $\mathcal{C}_\infty$ 50
- 2.5 Examples 71

3 Triangulated persistence Fukaya categories 97
- 3.1 Main symplectic topology applications 97
- 3.2 Filtrations in Floer homology and Fukaya categories 103
- 3.3 Genuinely filtered Fukaya categories 108
- 3.4 Proofs of the main symplectic applications 153
- 3.5 The geometry behind TPCs 168

References 189

Chapter 1

Introduction

The last 40 years have seen spectacular advances in symplectic topology. Most of them, especially those exhibiting aspects of symplectic rigidity, exploit algebraic structures that encode the behavior of moduli spaces of solutions of Cauchy–Riemann-type equations associated to (variants of) the symplectic action functional. Typical examples of such structures are Floer homology, Gromov–Witten invariants, and the Fukaya category. Given that these structures are, in essence, associated to a functional, they can be expected to admit refinements endowed with a finer structure reflective of an underlying filtration. Making this statement precise and incorporating this filtration in the respective algebraic structures is sometimes straightforward from the algebraic viewpoint. For instance, in favorable cases, the Floer complex – just like its more down-to-earth precursor, the Morse complex – is filtered, and its homology is a persistence module.

In other cases, such as that of the derived Fukaya category, which is the one that interests us here, this is far from immediate. In this memoir, we set up a new algebraic structure, called a *triangulated persistence category* (TPC), precisely to deal with this situation. This structure puts together persistence and triangulation and is a refinement of the notion of triangulated category. The construction is abstract and applicable to a variety of contexts unrelated to symplectic topology, as explained in more detail below, in Section 1.1.

The derived Fukaya category has a triangulated structure, and we show that, under certain constraints, it does admit a TPC refinement that is unique up to equivalence. The construction of this refinement and its uniqueness are delicate and require some novel geometric and algebraic steps. We describe in more detail the results and the constructions involved in Section 1.2. These constructions are also of independent interest.

A natural application of the construction of Fukaya-type triangulated persistence categories is a rigidity result for spaces of Lagrangian submanifolds. To fix ideas, let (X, ω) be a symplectic manifold. It is well known since the pioneering work of Gromov and Floer that closed Lagrangian submanifolds $L \subset X$ subject to certain *purely topological* constraints – the one used in this memoir is exactness – exhibit strong, and often surprising, rigidity properties that are intrinsically symplectic. Generally, this form of rigidity reflects individual properties of each of the Lagrangians in the fixed class. Two famous examples that have structured much of the modern work in the subject are the Arnold conjecture and the nearby Lagrangian conjecture, also due to Arnold.

In this work, we show that, in the same setting, a global form of rigidity is in effect. More precisely, let (X, ω) be a Liouville manifold that satisfies an algebraic finiteness condition that will be made explicit below. The set of closed, exact Lagrangians in X is endowed with a class of metrics, called *symplectic fragmentation* metrics, with some remarkable properties (see Corollary 3.7 for a more precise version and details):

- Up to a multiplicative constant, these metrics are dominated by the spectral metric (which is itself bounded above by the Hofer metric); thus, they carry symplectic content.
- The non-degeneracy of these metrics can be viewed as a form of Gromov's non-squeezing theorem, in the sense that the distance between two Lagrangians has a lower bound that can be expressed in terms of a purely geometric quantity: the supremum of the radii of standard symplectic balls that embed in a certain position relative to the two Lagrangians.
- The metrics are finite and thus they allow meaningful comparison of Lagrangians that are very different as smooth submanifolds (non-isotopic, or of different homotopy types), when classical metrics, such as the Hofer distance, are infinite.
- At the same time, they also satisfy a stability property for intersections, in the sense that, given two transverse Lagrangians L and N, if a third Lagrangian L' is sufficiently close to L in one of these metrics, then the number of intersections of L' with N cannot be smaller than the number of intersections of N with L.

The relation between this statement and the notion of TPC is that if a triangulated category admits a TPC refinement, then, by the main algebraic result in this memoir, its exact triangles are endowed with a so-called *persistence* triangular weight. The set of objects of a triangulated category, endowed with such a triangular weight, is easily seen to carry a family of natural pseudometrics called *fragmentation* pseudometrics. The symplectic fragmentation metric mentioned above is deduced from the fragmentation pseudometrics associated to the Fukaya TPC.

Remark 1.1. Precursors of the metrics introduced here have appeared in [10], based on Lagrangian cobordism machinery. However, the constructions in that paper lacked the proper algebraic setting, with the consequence that the finiteness of the distance between two Lagrangians depended on the existence of certain Lagrangian cobordisms. This issue was addressed, in part, in [9] through considerations involving immersed Lagrangians, which allow the construction of an abundance of immersed cobordisms. However, the immersed cobordism approach is technically very delicate, and less natural than the one proposed here, with the consequence that it leads to a family of metrics that are extremely hard to estimate.

1.1 Persistence and triangulation

Persistence theory, introduced in several pioneering works [14, 18, 23, 29, 33, 42, 62, 64], is an abstract framework that emerged from investigations in parts of data science as well as in topology, formalizing the structure and properties of a class of phenomena that are most easily seen in the homology of a chain complex (C, d) endowed with an increasing filtration $(C^{\leq\alpha}, d) \subset (C, d)$ of subcomplexes parametrized by $\alpha \in \mathbb{R}$. The homologies of the subcomplexes form a family $\{H(C^{\leq\alpha})\}_{\alpha\in\mathbb{R}}$ whose members are related by maps $i_{\alpha,\beta} : H(C^{\leq\alpha}) \to H(C^{\leq\beta})$, $\alpha \leq \beta$, subject to obvious compatibilities. This is an example of a persistence module. Given two filtered complexes (C, d) and (D, d) that are quasi-isomorphic, it is possible to compare them by the so-called interleaving distance. Its definition is based on the fact that the space of linear maps $v : C \to D$ is itself filtered by the "shift" of a map: v is of shift at most r if $v(C^{\leq\alpha}) \subset D^{\leq\alpha+r}$ for all $\alpha \in \mathbb{R}$. Using this, given two chain maps $\phi : C \to D$, $\psi : D \to C$ such that $\psi \circ \phi$ is chain homotopic to $\mathbb{1}_C$, there is a natural measurement for how far the composition $\psi \circ \phi$ is from the identity, namely the infimum of the "shifts" of chain homotopies $h : C \to C$ such that $dh + hd = \psi \circ \phi - \mathbb{1}_C$. The machinery of persistence modules is much more developed than the few elements mentioned here. For a survey of this topic and its applications in various branches of mathematics, see the monographs and papers by Edelsbrunner [28], Oudot [46], Chazal–de Silva–Glisse–Oudot [17], Polterovich–Shelukhin [48], Polterovich–Rosen–Samvelyan–Zhang [47], and Kislev–Shelukhin [41]. In particular, there is a beautiful interpretation of the bottleneck distance in terms of so-called barcodes [5, 59].

The main question that we address in the algebraic part of this memoir is independent of symplectic considerations:

How can one use a persistence-type structure on the morphisms of a category to compare not only (quasi-)isomorphic objects but rather define a pseudometric on the set of all objects?

We provide here a solution to this question based on mixing persistence with triangulation understood in the sense of triangulated categories as introduced by Puppe [50] and Verdier [60] in the early 1960s. Given a triangulated category $\mathcal{D}$, there is a simple notion of *triangular weight* w on $\mathcal{D}$ that we introduce in Section 2.1. This associates to each exact triangle Δ in $\mathcal{D}$ a non-negative number $w(\Delta)$ satisfying a couple of properties. The most relevant of them is a weighted variant of the octahedral axiom (we will give a more precise definition later). A basic example of a triangular weight is the flat one: it associates to each exact triangle the value 1. The interest of triangular weights is that they naturally lead to *fragmentation pseudometrics* on $\mathrm{Obj}(\mathcal{D})$ (we assume here that $\mathcal{D}$ is small) defined roughly as follows (see Section 2.1 for details). Such a pseudometric depends on a family of objects $\mathcal{F}$ of $\mathcal{D}$. With $\mathcal{F}$ fixed, and up to a certain normalization, the pseudodistance $d^{\mathcal{F}}(X, Y)$ between $X, Y \in \mathrm{Obj}(\mathcal{D})$ is (the symmetrization of) the infimum of the total weight of

exact triangles needed to construct iteratively X out of Y by only attaching cones over morphisms with domain in $\mathcal{F}$. The weighted octahedral axiom implies that this $d^{\mathcal{F}}$ satisfies the triangle inequality. Using such pseudometrics one can analyze rigidity properties of various categories by exploring the induced topology on $\mathrm{Obj}(\mathcal{D})$.

The main algebraic part of the memoir is contained in Chapter 2, and its aim is to use persistence machinery to produce certain non-flat triangular weights. The main tool, as already mentioned above, is a refinement of triangulated categories, called *triangulated persistence categories* (TPCs). A triangulated persistence category has two main properties. First, it is a persistence category, a natural notion we introduce in Section 2.2. This is a category $\mathcal{C}$ whose morphisms $\hom_{\mathcal{C}}(A, B)$ are persistence modules, and whose composition is compatible with the persistence structure. (See [47] for a general introduction to persistence module theory.) The second main structural property of TPCs is that the objects of $\mathcal{C}$, together with the 0-persistence level morphisms $\hom^0_{\mathcal{C}}(A, B)$, have the structure of a triangulated category $\mathcal{C}_0$. The formal definition of TPCs is given in Section 2.3.

It is natural to associate to a persistence category $\mathcal{C}$ a limit category $\mathcal{C}_\infty$ that has the same objects as $\mathcal{C}$ and has as morphisms the ∞-limits of the morphisms in $\mathcal{C}$. In a different direction, a natural notion in a persistence category is that of an r-acyclic object: K is called *r-acyclic* if its identity morphism $\mathbb{1}_K \in \hom^0_{\mathcal{C}}(K, K)$ is 0 in $\hom^r_{\mathcal{C}}(K, K)$. The acyclic objects for all $r \geq 0$ form a full subcategory $\mathcal{AC}$ of $\mathcal{C}$ that is also a persistence category, and in case $\mathcal{C}$ is a TPC, it is easy to see that $\mathcal{AC}$ is also a TPC. In particular, $\mathcal{AC}_0$ is triangulated.

These notions are tied together by the classical construction of Verdier localization. Indeed, assuming as above that $\mathcal{C}$ is a TPC, we will see that $\mathcal{C}_\infty$ coincides with the Verdier localization of $\mathcal{C}_0$ with respect to $\mathcal{AC}_0$. In particular, the category $\mathcal{C}_\infty$ is also triangulated.

We can now state the main result of the algebraic part of the memoir (restated more precisely in Theorem 2.65).

Theorem A. *If $\mathcal{C}$ is a triangulated persistence category, then, with the notation above, the Verdier localization $\mathcal{C}_\infty$ admits a non-flat triangular weight induced by the persistence structure of $\mathcal{C}$.*

The construction of this triangular weight is based on a definition of a class of weighted triangles in the category $\mathcal{C}$ itself. With this definition, the exact triangles in $\mathcal{C}_0$ have weight 0, but there are also other triangles in $\mathcal{C}$ of arbitrary positive weight. While the category $\mathcal{C}$, together with the class of finite-weight triangles, is not triangulated – even the formal expression of these triangles in $\mathcal{C}$ does not fit the axioms of triangulated categories – the properties of these triangles are sufficient to induce a triangular weight on the exact triangles of $\mathcal{C}_\infty$.

In summary, if a triangulated category $\mathcal{D}$ admits a TPC refinement – that is, a TPC $\mathcal{C}$ such that $\mathcal{C}_\infty = \mathcal{D}$ (as triangulated categories) – then $\mathcal{D}$ carries a non-flat

triangular weight induced by the persistence structure of $\mathcal{C}$. As a result, this construction provides a technique to build non-discrete fragmentation pseudometrics on the objects of $\mathcal{D}$.

Some classes of examples are discussed in Section 2.5. Triangulated persistence categories are expected to be of use beyond the field of symplectic topology, and Chapter 2, which is essentially self-contained, can be read independently of the symplectic considerations that appear in Chapter 3.

Remark 1.2. Even the fragmentation pseudometrics associated to the flat weight are of interest. Many qualitative questions concerning numerical lower bounds for the complexity of certain geometric objects can be understood by means of inequalities involving such fragmentation pseudometrics. Classical examples are the Morse inequalities, the Lusternik–Schnirelmann inequality, as well as, in symplectic topology, the inequalities predicted by the Arnold conjectures. Remarkable results based on measurements using this flat weight and applied to the study of endofunctors have appeared recently in work of Orlov [45], as well as Dimitrov–Haiden–Katzarkov–Kontsevich [26] and Fan–Filip [30].

1.2 TPC refinements of the Fukaya derived category

Here is an overview of the geometric part of the memoir (Chapter 3). The main step here is to consider a finite family $\mathcal{X}$ of closed, exact Lagrangians in a symplectic manifold X, assumed in general position, and construct a TPC refinement of the derived Fukaya category of $\mathcal{X}$.

There are quite a few nuances here. First, this requires the construction of a filtered Fukaya-type category with objects the elements of $\mathcal{X}$, endowed with all possible primitives. A *weakly* filtered such category has been constructed in [10], but obtaining a genuinely filtered A_∞-structure is more delicate. It requires careful control of energy estimates (and the technique we use restricts us to finite families $\mathcal{X}$), but also the use of "cluster"-type moduli spaces, that mix J-holomorphic polygons and Morse trajectories. Fortunately, such moduli spaces have been studied and used frequently since [21], for instance in [15, 16].

The resulting filtered Fukaya category $\mathcal{F}uk(\mathcal{X})$ depends, of course, on choices of auxiliary structures such as perturbation data that we omit from the notation here. The next step is to pursue the construction of the derived version. As in the non-filtered version, this part is purely algebraic and applies to any filtered A_∞-category. Nonetheless, there are some significant differences with respect to the non-filtered case. Uniqueness up to equivalence is considerably more delicate to achieve because several basic algebraic A_∞-tools, such as the Hochschild complex and related constructions, require significant adjustment to adapt to the filtered setting. Moreover,

at a more conceptual level, the two natural constructions of the derived category – one based on filtered twisted complexes and the other on the Yoneda embedding and A_∞ filtered modules – both lead to useful natural notions, but not to equivalent ones. Denote by $\mathcal{CF}uk(\mathcal{X})$ the version based on filtered modules. Let $D\mathcal{F}uk(\mathcal{X})$ be the usual, unfiltered, derived Fukaya category of $\mathcal{X}$, and assume that $\mathcal{F} \subset \mathcal{X}$ is a family of triangular generators for $D\mathcal{F}uk(\mathcal{X})$. Fix also a second such family $\mathcal{F}'$ with each element being a small generic Hamiltonian deformation of a corresponding element in $\mathcal{F}$.

The main statement is the following – again in simplified form (the full statement is in Theorem 3.4):

Theorem B. *The category $\mathcal{CF}uk(\mathcal{X})$ is a TPC and it is independent of the defining data up to TPC equivalence. Moreover, $\mathcal{CF}uk(\mathcal{X})_\infty$ is triangulated equivalent to $D\mathcal{F}uk(\mathcal{X})$. Finally, there exists a fragmentation metric on $\mathcal{X}$ that is independent of the choices used in the construction of $\mathcal{CF}uk(\mathcal{X})$ and is defined by*

$$D^{\mathcal{F},\mathcal{F}'} = \max\{D^{\mathcal{F}}, D^{\mathcal{F}'}\},$$

where $D^{\mathcal{F}}$ are the shift-invariant versions of the fragmentation pseudometrics $d^{\mathcal{F}}$ constructed as outlined in Section 1.1.

One delicate point worth emphasizing here is that while we expect $\mathcal{CF}uk(\mathcal{X})$ to be unique up to *canonical* equivalence, the machinery in this memoir does not produce fully canonical equivalences (see Theorem 3.12).

Of course, as the set $\mathcal{X}$ is finite, this metric $D^{\mathcal{F},\mathcal{F}'}$ might appear to be uninteresting; however, the more precise result (Theorem 3.4) shows that the pseudometrics $D^{\mathcal{F}}$ satisfy some remarkable properties (see also Remark 3.5). These properties are then used to analyze how the pseudometrics change when the family $\mathcal{X}$ increases. Ultimately, this leads to the definition of the metric on the space of all closed exact Lagrangians that was claimed earlier in the introduction. This is stated more precisely in Corollary 3.7.

The construction of TPCs is inspired by recent constructions in symplectic topology and, in particular, by the shadow pseudometrics introduced in [10] and [9] in the study of Lagrangian cobordism. This aspect is discussed in Section 3.5.1. The construction of the filtered Fukaya category and the associated TPC are expected to be of independent interest.

Chapter 2

Triangulation persistence categories: Algebra 101

This chapter contains the main algebraic machinery introduced in the memoir and is self-contained, except for some basic elements of homological algebra, as can be found in [61].[1]

In Section 2.1 we briefly introduce the notion of triangular weight and discuss its application to measuring the complexity of cone decompositions in triangulated categories. In Section 2.2 we introduce persistence categories, which are, in short, categories enriched by persistence modules. Triangulated persistence categories are introduced in Section 2.3. In Section 2.4 we prove the main algebraic result of the chapter, namely that the ∞-level of a TPC carries a specific triangular weight induced by the persistence structure. Finally, in Section 2.5 we discuss some classes of natural TPC examples that are not symplectic in nature (the symplectic examples are deferred to Chapter 3).

2.1 Triangular weights

In this subsection we introduce triangular weights associated to a triangulated category $\mathcal{D}$. Using such a triangular weight w on $\mathcal{D}$ we define a class of so-called fragmentation pseudometrics $d_w^{\mathcal{F}}$ on $\mathrm{Obj}(\mathcal{D})$. All categories used in this memoir (in particular, $\mathcal{D}$) are assumed to be small unless otherwise indicated.

Definition 2.1. Let $\mathcal{D}$ be a triangulated category and denote by $\mathcal{T}_{\mathcal{D}}$ its class of exact triangles. A *triangular weight* w on $\mathcal{D}$ is a function

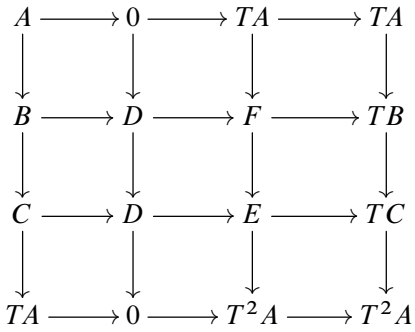
$$w : \mathcal{T}_{\mathcal{D}} \to [0, \infty)$$

that satisfies properties (i) and (ii) below:

(i) [Weighted octahedral axiom] Assume that the triangles $\Delta_1 : A \to B \to C \to TA$ and $\Delta_2 : C \to D \to E \to TC$ are both exact. There are exact triangles $\Delta_3 : B \to D \to F \to TB$ and $\Delta_4 : TA \to F \to E \to T^2A$ making the diagram below

[1] A version of this chapter appeared earlier as an independent preprint [11]. The only changes compared to [11], besides minor corrections of imprecisions, concern the relations to Verdier localization in Section 2.3.2.

commute, except for the bottom-right square, which anti-commutes:



and such that

$$w(\Delta_3) + w(\Delta_4) \leq w(\Delta_1) + w(\Delta_2).$$

(ii) [Normalization] There exists some $w_0 \in [0, \infty)$ such that $w(\Delta) \geq w_0$ for all $\Delta \in \mathcal{T}_{\mathcal{D}}$ and $w(\Delta') = w_0$ for all triangles Δ' of the form $0 \to X \xrightarrow{\mathbb{1}_X} X \to 0$, $X \in \mathrm{Obj}(\mathcal{D})$, and their rotations. Moreover, in the diagram in (i), if $B = 0$ we may take Δ_3 to be

$$\Delta_3 : 0 \to D \to D \to 0. \tag{2.1}$$

Remark 2.2. (a) Neglecting the weight constraints, given the triangles Δ_1, Δ_2, Δ_3 as in point (i), the octahedral axiom is easily seen to imply the existence of Δ_4 making the diagram commutative, as in the definition.

(b) The condition in point (ii), above equation (2.1), can be reformulated as a replacement property for exact triangles in the following sense: if $\Delta_2 : C \to D \to E \to TC$ is exact and C is isomorphic to A' $(= TA)$, then there is an exact triangle $A' \to D \to E \to TA'$ of weight at most $w(\Delta_2) + w(\Delta_1) - w_0$, where Δ_1 is the exact triangle $T^{-1}A' \to 0 \to C \to A'$.

Given an exact triangle $\Delta : A \to B \xrightarrow{f} C \to TA$ in $\mathcal{D}$ and any $X \in \mathrm{Obj}(\mathcal{D})$, there is an associated exact triangle $X \oplus \Delta : A \to X \oplus B \xrightarrow{\mathbb{1}_X \oplus f} X \oplus C \to TA$ and a similar one, $\Delta \oplus X$. We say that a triangular weight w on $\mathcal{D}$ is *subadditive* if, for any exact triangle $\Delta \in \mathcal{T}_{\mathcal{D}}$ and any object X of $\mathcal{D}$, we have

$$w(X \oplus \Delta) \leq w(\Delta)$$

and similarly for $\Delta \oplus X$.

The simplest example of a triangular weight on a triangulated category $\mathcal{D}$ is the flat one, $w_{fl}(\Delta) = 1$, for all triangles $\Delta \in \mathcal{T}_{\mathcal{D}}$. This weight is obviously subadditive. A weight that is not proportional to the flat one is called *non-flat*.

The interest of triangular weights comes from the next definition, which provides a measure of the complexity of cone decompositions in $\mathcal{D}$; this leads in turn to the definition of corresponding pseudometrics on the set $\mathrm{Obj}(\mathcal{D})$.

Definition 2.3. Fix a triangulated category $\mathcal{D}$ together with a triangular weight w on $\mathcal{D}$. Let X be an object of $\mathcal{D}$. An *iterated cone decomposition* D of X with *linearization* $\ell(D) = (X_1, X_2, \ldots, X_n)$ consists of a family of exact triangles in $\mathcal{D}$:

$$\begin{cases} \Delta_1 : X_1 \to 0 \to Y_1 \to TX_1 \\ \Delta_2 : X_2 \to Y_1 \to Y_2 \to TX_2 \\ \Delta_3 : X_3 \to Y_2 \to Y_3 \to TX_3 \\ \qquad\qquad \vdots \\ \Delta_n : X_n \to Y_{n-1} \to X \to TX_n. \end{cases}$$

To accommodate the case $n = 1$ we set $Y_0 = 0$. The *weight* of such a cone decomposition is defined by

$$w(D) = \sum_{i=1}^{n} w(\Delta_i) - w_0.$$

This weight of cone decompositions naturally leads to a class of pseudometrics on the objects of $\mathcal{D}$, as follows.

Let $\mathcal{F} \subset \mathrm{Obj}(\mathcal{D})$. For two objects X, X' of $\mathcal{D}$, define

$$\begin{aligned} \delta^{\mathcal{F}}(X, X') = \inf\big\{ & w(D) \mid D \text{ is an iterated cone decomposition of } X \\ & \text{with linearization } (F_1, \ldots, T^{-1}X', \ldots, F_k), \\ & \text{where } F_i \in \mathcal{F}, k \geq 0 \big\}. \end{aligned} \tag{2.2}$$

Note that we allow here $k = 0$, i.e., the linearization of D is allowed to consist of only one element, $T^{-1}X'$, without using any elements F_i from the family $\mathcal{F}$. Fragmentation pseudometrics are obtained by symmetrizing $\delta^{\mathcal{F}}$, as below.

Proposition 2.4. *Let $\mathcal{D}$ be a triangulated category and let w be a triangular weight on $\mathcal{D}$. Fix $\mathcal{F} \subset \mathrm{Obj}(\mathcal{C})$ and define*

$$d^{\mathcal{F}} : \mathrm{Obj}(\mathcal{C}) \times \mathrm{Obj}(\mathcal{C}) \to [0, \infty) \cup \{+\infty\}$$

by

$$d^{\mathcal{F}}(X, X') = \max\{\delta^{\mathcal{F}}(X, X'), \delta^{\mathcal{F}}(X', X)\}.$$

(i) *The map $d^{\mathcal{F}}$ is a pseudometric, called* the fragmentation pseudometric *associated to w and $\mathcal{F}$.*

(ii) *If w is subadditive, then*

$$d^{\mathcal{F}}(A \oplus B, A' \oplus B') \leq d^{\mathcal{F}}(A, A') + d^{\mathcal{F}}(B, B') + w_0. \tag{2.3}$$

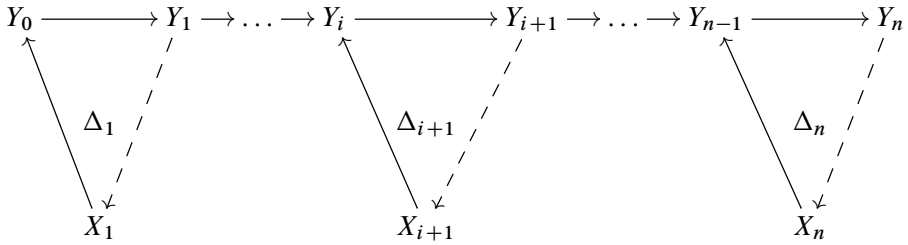
In particular, if $w_0 = 0$, then $\mathrm{Obj}(\mathcal{D})$ *with the operation given by $\oplus$ and the topology induced by $d^{\mathcal{F}}$ is an H-space. (Recall that a topological space*

is called an H-space *if there exists a continuous map* $\mu : X \times X \to X$ *with an identity element* e *such that* $\mu(e, x) = \mu(x, e) = x$ *for any* $x \in X$.)

The proof of Proposition 2.4 is based on simple manipulations with exact triangles. We will prove a similar statement in Section 2.3.4 in a more complicated setting and we will then briefly discuss in Section 2.3.4.1 how the arguments given in that case also imply Proposition 2.4.

Remark 2.5. (a) If $\mathcal{F}$ is invariant under translation, in the sense that $T\mathcal{F} \subset \mathcal{F}$ and, moreover, $\mathcal{F}$ is a family of triangular generators for $\mathcal{C}$, then the metric $d^{\mathcal{F}}$ admits finite values. This is not difficult to show by first proving that $\delta^{\mathcal{F}}(0, X)$ is finite for all $X \in \mathrm{Obj}(\mathcal{C})$ (it is immediate that $\delta^{\mathcal{F}}(X, 0)$ is finite).

(b) It is sometimes useful to view an iterated cone decomposition as in Definition 2.3 as a sequence of objects and maps forming the successive triangles Δ_i below



(2.4)

where the dotted arrows represent maps $Y_i \to TX_i$ and $Y_0 = 0$, $Y_n = X$.

(c) The definition of fragmentation pseudometrics is quite flexible and there are a number of possible variants. One of them will be useful later. Instead of $\delta^{\mathcal{F}}$ as given in (2.2), we may use

$$\underline{\delta}^{\mathcal{F}}(X, X') = \inf\Big\{\sum_{i=1}^{n} w(\Delta_i) \mid \Delta_i \text{ are successive exact triangles as in (2.4)} \\ \text{with } Y_1 = X',\ X = Y_n, \text{ and } X_i \in \mathcal{F},\ n \in \mathbb{N}\Big\}. \tag{2.5}$$

For this to be coherent we need to assume here $0 \in \mathcal{F}$. Comparing with the definition of $\delta^{\mathcal{F}}$ in (2.2), $\underline{\delta}^{\mathcal{F}}$ corresponds to only taking into account cone decompositions with linearization $(T^{-1}X', F_1, \ldots, F_n)$ and with the first triangle $\Delta_1 : T^{-1}X' \to 0 \to X' \to X'$. There are two advantages to this expression: the first is that it is trivial to see in this case that $\underline{\delta}^{\mathcal{F}}$ satisfies the triangle inequality, which does not even require the weighted octahedral axiom. The other advantage is that one starts the sequence of triangles from X', and thus the negative translate $T^{-1}X'$ is not needed to define $\underline{\delta}^{\mathcal{F}}$. There is an associated fragmentation pseudometric $\underline{d}^{\mathcal{F}}$ obtained by symmetrizing $\underline{\delta}^{\mathcal{F}}$, and this satisfies a formula similar to (2.3). Of course, the disadvantage of this fragmentation pseudometric is that it is larger than $d^{\mathcal{F}}$ and thus more often infinite.

2.2 Persistence categories

We introduce in this section the notion of a persistence category – a category whose morphisms are persistence modules and such that composition respects the persistence structure – and then proceed with a number of related structures and immediate properties.

2.2.1 Basic definitions

View the real axis $\mathbb{R}$ as a category with $\mathrm{Obj}(\mathbb{R}) = \{x \mid x \in \mathbb{R}\}$ and, for any $x, y \in \mathrm{Obj}(\mathbb{R})$, the hom-set

$$\hom_{\mathbb{R}}(x, y) = \begin{cases} i_{x,y} & \text{if } x \leq y, \\ \emptyset & \text{if } x > y. \end{cases}$$

By definition, for any $x \leq y \leq z$ in $\mathbb{R}$, $i_{y,z} \circ i_{x,y} = i_{x,z}$. We denote this category by $(\mathbb{R}, \leq)$. It admits an additive structure. Explicitly, consider the bifunctor $\oplus : (\mathbb{R}, \leq) \times (\mathbb{R}, \leq) \to (\mathbb{R}, \leq)$ defined by $\oplus(r, s) := r + s$, where $0 \in \mathbb{R}$ is the zero object and, for any two pairs $(r, s), (r', s') \in \mathrm{Obj}((\mathbb{R}, \leq) \times (\mathbb{R}, \leq))$,

$$\hom_{(\mathbb{R},\leq)\times(\mathbb{R},\leq)}((r, s), (r', s')) = \begin{cases} (i_{r,r'}, i_{s,s'}) & \text{if } r \leq r' \text{ and } s \leq s', \\ \emptyset & \text{otherwise,} \end{cases}$$

and further $\oplus(i_{r,r'}, i_{s,s'}) := i_{r+s,r'+s'} \in \hom_{(\mathbb{R},\leq)}(r + s, r' + s')$. Fix a ground field $\mathbf{k}$ and denote by $\mathrm{Vect}_{\mathbf{k}}$ the category of $\mathbf{k}$-vector spaces.

Definition 2.6. A category $\mathcal{C}$ is called a *persistence category* if it is endowed with the following additional structure. For any $A, B \in \mathrm{Obj}(\mathcal{C})$, we are given a functor $E_{A,B} : (\mathbb{R}, \leq) \to \mathrm{Vect}_{\mathbf{k}}$ such that the following two conditions are satisfied:

(i) The hom-set in $\mathcal{C}$ is $\hom_{\mathcal{C}}(A, B) = \{(f, r) \mid f \in E_{A,B}(r)\}$. We write $\hom^r_{\mathcal{C}}(A, B) := E_{A,B}(r)$, or simply $\hom^r(A, B)$, when the ambient category $\mathcal{C}$ is not emphasized.

(ii) The composition $\circ : \hom^r_{\mathcal{C}}(A, B) \times \hom^s_{\mathcal{C}}(B, C) \to \hom^{r+s}_{\mathcal{C}}(A, C)$ in $\mathcal{C}$ is a natural transformation from $E_{A,B} \times E_{B,C}$ to $E_{A,C} \circ \oplus$ (with $\oplus$ the product $(\mathbb{R}, \leq) \times (\mathbb{R}, \leq) \to (\mathbb{R}, \leq)$). Explicitly, the following diagram commutes:

$$\begin{array}{ccc} \hom^r_{\mathcal{C}}(A, B) \times \hom^s_{\mathcal{C}}(B, C) & \xrightarrow{\circ_{(r,s)}} & \hom^{r+s}_{\mathcal{C}}(A, C) \\ \Big\downarrow{\scriptstyle E_{A,B}(i_{r,r'})\times E_{B,C}(i_{s,s'})} & & \Big\downarrow{\scriptstyle E_{A,C}(i_{r+s,r'+s'})} \\ \hom^{r'}_{\mathcal{C}}(A, B) \times \hom^{s'}_{\mathcal{C}}(B, C) & \xrightarrow{\circ_{(r',s')}} & \hom^{r'+s'}_{\mathcal{C}}(A, C) \end{array}$$

Remark 2.7. Item (i) means that each hom-set $\hom_{\mathcal{C}}(A, B)$ is a persistence $\mathbf{k}$-module with persistence structure morphisms $E_{A,B}(i_{r,s})$ for any $r \leq s$ in $\mathbb{R}$. Here, we use the weakest possible definition of a persistence $\mathbf{k}$-module, in the sense that no regularities, such as the finiteness of the dimension of $\hom^r_{\mathcal{C}}(A, B)$ or the semi-continuity when changing the parameter r, are required (see [47, Subsection 1.1]).

We will often denote an element in $\hom_{\mathcal{C}}(A, B)$ by a single symbol $\bar{f}$ instead of a pair (f, r). We will use the notation $\lceil \bar{f} \rceil = r$ to denote the real number r and refer to this number as the *shift* (or persistence level) of $\bar{f}$. For each $A \in \mathrm{Obj}(\mathcal{C})$, the identity $\bar{\mathbb{1}}_A := (\mathbb{1}_A, 0) \in \hom^0_{\mathcal{C}}(A, A)$ is of shift 0. If one of the objects A or B is the zero object, then $\hom_{\mathcal{C}}(A, B)$ contains only the zero morphism, denoted by 0, and it lies in $\hom^r_{\mathcal{C}}(A, B)$ for any $r \in \mathbb{R}$. For brevity, we will denote from now on the structural morphisms $E_{A,B}(i_{r,s})$ by $i_{r,s}$.

A persistence structure allows us to consider morphisms that are identified up to r-shift and, similarly, objects that are negligible up to a shift by r.

Definition 2.8. Fix a persistence category $\mathcal{C}$.

(i) For $f, g \in \hom^{\alpha}_{\mathcal{C}}(A, B)$, we say that f and g are *r-equivalent* for some $r \geq 0$ if
$$i_{\alpha,\alpha+r}(f - g) = 0.$$
We write $f \simeq_r g$ if f and g are r-equivalent.

(ii) Two morphisms $f \in \hom^{\alpha}_{\mathcal{C}}(A, B)$ and $g \in \hom^{\beta}_{\mathcal{C}}(A, B)$ are *∞-equivalent*, written $f \simeq_\infty g$, if there exist $r, r' \geq 0$ with $\alpha + r = \beta + r'$ such that $i_{\alpha,\alpha+r}(f) = i_{\beta,\beta+r'}(g)$.

(iii) An object $K \in \mathrm{Obj}(\mathcal{C})$ is called *r-acyclic* for some $r \geq 0$ if its identity morphism $\mathbb{1}_K \in \hom^0_{\mathcal{C}}(K, K)$ has the property that $\mathbb{1}_K \simeq_r 0$.

Obviously, if $f \simeq_r g$ then $f \simeq_s g$ for all $s \geq r$. Notice also that $\simeq_r$ is indeed an equivalence relation. Indeed, for $r \neq \infty$ this follows immediately from the fact that $i_{\alpha,\beta} : \hom^{\alpha}_{\mathcal{C}}(A, B) \to \hom^{\beta}_{\mathcal{C}}(A, B)$ is a linear map, and it is an easy exercise for $r = \infty$.

Definition 2.9. Given a persistence category $\mathcal{C}$, there are two categories naturally associated to it, as follows:

(i) the 0-*level* of $\mathcal{C}$, denoted by $\mathcal{C}_0$, which is the category with the same objects as $\mathcal{C}$ and, for any $A, B \in \mathrm{Obj}(\mathcal{C})$, with $\hom_{\mathcal{C}_0}(A, B) := \hom^0_{\mathcal{C}}(A, B)$.

(ii) the *limit category* (or ∞-*level*) of $\mathcal{C}$, denoted by $\mathcal{C}_\infty$, which again has the same objects as $\mathcal{C}$ but, for any $A, B \in \mathrm{Obj}(\mathcal{C})$,
$$\hom_{\mathcal{C}_\infty}(A, B) := \varinjlim_{\alpha \to \infty} \hom^{\alpha}_{\mathcal{C}}(A, B),$$

where the direct limit is taken with respect to the morphisms

$$i_{\alpha,\beta} : \hom_{\mathcal{C}}^{\alpha}(A, B) \to \hom_{\mathcal{C}}^{\beta}(A, B)$$

for any $\alpha \leq \beta$.

Remark 2.10. (a) In general, a persistence category is not pre-additive as the hom-sets $\hom_{\mathcal{C}}(A, B)$ are generally not abelian groups. However, it is easy to see that both $\mathcal{C}_0$ and $\mathcal{C}_\infty$ are pre-additive (the proof is immediate in the first case and a simple exercise in the second).

(b) The limit category $\mathcal{C}_\infty$ can be equivalently defined as the quotient category $\mathcal{C}/\simeq_\infty$, which is defined by

$$\mathrm{Obj}(\mathcal{C}/\simeq_\infty) = \mathrm{Obj}(\mathcal{C}) \quad \text{and} \quad \hom_{\mathcal{C}/\simeq_\infty}(A, B) = \hom_{\mathcal{C}}(A, B)/\simeq_\infty .$$

Two objects $A, B \in \mathrm{Obj}(\mathcal{C})$ are said to be 0-*isomorphic*, and we write $A \equiv B$, if they are isomorphic in the category $\mathcal{C}_0$. This is obviously an equivalence relation, and it preserves r-acyclics in the sense that if $K \simeq_r 0$ and $K \equiv K'$, then $K' \simeq_r 0$.

2.2.2 Persistence functors

Persistence categories come with associated notions of persistence functors and natural transformations relating them, as described below.

Definition 2.11. Given two persistence categories $\mathcal{C}$ and $\mathcal{C}'$, a *persistence functor* $\mathcal{F} : \mathcal{C} \to \mathcal{C}'$ is a functor that is compatible with the persistence structures. More explicitly, the action of $\mathcal{F}$ on morphisms restricts to maps $(\mathcal{F}_{A,B})_r : \hom_{\mathcal{C}}^{r}(A, B) \to \hom_{\mathcal{C}'}^{r}(\mathcal{F}(A), \mathcal{F}(B))$ defined for any $A, B \in \mathrm{Obj}(\mathcal{C})$ and $r \in \mathbb{R}$. Moreover, for every $r \leq s$ we have the following commutative diagram:

$$\begin{array}{ccc} \hom_{\mathcal{C}}^{r}(A, B) & \xrightarrow{(\mathcal{F}_{A,B})_r} & \hom_{\mathcal{C}'}^{r}(\mathcal{F}(A), \mathcal{F}(B)) \\ \Big\downarrow{\scriptstyle i_{r,s}^{\mathcal{C}}} & & \Big\downarrow{\scriptstyle i_{r,s}^{\mathcal{C}'}} \\ \hom_{\mathcal{C}}^{s}(A, B) & \xrightarrow{(\mathcal{F}_{A,B})_s} & \hom_{\mathcal{C}'}^{s}(\mathcal{F}(A), \mathcal{F}(B)) \end{array}$$

where $i_{r,s}^{\mathcal{C}}$ and $i_{r,s}^{\mathcal{C}'}$ are persistence structure maps in $\mathcal{C}$ and $\mathcal{C}'$, respectively. In particular, for each $\bar{f} \in \hom_{\mathcal{C}}(A, B)$ with $\lceil \bar{f} \rceil = r$, we have $\lceil \mathcal{F}_{A,B}(\bar{f}) \rceil = r$.

For any functor $E : (\mathbb{R}, \leq) \to \mathrm{Vect}_{\mathbf{k}}$ and $\alpha \in \mathbb{R}$, we denote by $\Sigma^\alpha E : (\mathbb{R}, \leq) \to \mathrm{Vect}_{\mathbf{k}}$ the α-shift of E, defined by $\Sigma^\alpha E(r) = E(r + \alpha)$ and $\Sigma^\alpha E(i_{r,s}) = E(i_{r+\alpha,s+\alpha})$ for any $i_{r,s} : r \to s$, $r \leq s$.

Definition 2.12. Given two persistence functors between two persistence categories $\mathcal{F}, \mathcal{G} : \mathcal{C} \to \mathcal{C}'$, a *persistence natural transformation* $\eta : \mathcal{F} \to \mathcal{G}$ is a natural transformation for which there exists $r \in \mathbb{R}$ such that for any $A \in \mathrm{Obj}(\mathcal{C})$, the morphism

$\eta_A : \mathcal{F}(A) \to \mathcal{G}(A)$ belongs to $\hom^r_{\mathcal{C}'}(\mathcal{F}(A), \mathcal{G}(A))$. We say that η is a natural transformation of shift r.

Remark 2.13. (a) The morphisms $\eta_A : \mathcal{F}(A) \to \mathcal{G}(A)$, $A \in \mathrm{Obj}(\mathcal{C})$, give rise to the following commutative diagrams for all $X \in \mathrm{Obj}(\mathcal{C})$ and for any $\alpha \leq \beta \in \mathbb{R}$:

$$\begin{array}{ccc} \hom^{\alpha}_{\mathcal{C}'}(X, \mathcal{F}(A)) & \xrightarrow{\eta_A \circ} & \hom^{\alpha+r}_{\mathcal{C}'}(X, \mathcal{G}(A)) \\ \downarrow{\scriptstyle i_{\alpha,\beta}} & & \downarrow{\scriptstyle i_{\alpha+r,\beta+r}} \\ \hom^{\beta}_{\mathcal{C}'}(X, \mathcal{F}(A)) & \xrightarrow{\eta_A \circ} & \mathrm{Mor}^{\beta+r}_{\mathcal{C}'}(X, \mathcal{G}(A)) \end{array} \tag{2.6}$$

and

$$\begin{array}{ccc} \hom^{\alpha}_{\mathcal{C}'}(\mathcal{G}(A), X) & \xrightarrow{\circ \eta_A} & \hom^{\alpha+r}_{\mathcal{C}'}(\mathcal{F}(A), X) \\ \downarrow{\scriptstyle i_{\alpha,\beta}} & & \downarrow{\scriptstyle i_{\alpha+r,\beta+r}} \\ \hom^{\beta}_{\mathcal{C}'}(\mathcal{G}(A), X) & \xrightarrow{\circ \eta_A} & \hom^{\beta+r}_{\mathcal{C}'}(\mathcal{F}(A), X) \end{array} \tag{2.7}$$

(b) Given two persistence categories $\mathcal{C}, \mathcal{C}'$, the persistence functors themselves form a persistence category denoted by $\mathcal{P}\mathrm{Fun}(\mathcal{C}, \mathcal{C}')$, where

$$\hom_{\mathcal{P}\mathrm{Fun}(\mathcal{C},\mathcal{C}')}(\mathcal{F}, \mathcal{G}) = \big\{(\eta, r) \mid \eta \text{ is a natural transformation from } \mathcal{F} \text{ to } \mathcal{G} \text{ of shift } r\big\}.$$

When $\mathcal{C} = \mathcal{C}'$, simply denote $\mathcal{P}\mathrm{Fun}(\mathcal{C}, \mathcal{C}')$ by $\mathcal{P}\mathrm{End}(\mathcal{C})$. It is easy to verify that $\mathcal{P}\mathrm{End}(\mathcal{C})$ admits a strict monoidal structure.

Definition 2.14. Let $\mathcal{C}', \mathcal{C}''$ be two persistence categories. A persistence functor $\mathcal{F} : \mathcal{C}' \to \mathcal{C}''$ is called an *equivalence of persistence categories* (or *persistence equivalence*) if there exists a persistence functor $\mathcal{G} : \mathcal{C}'' \to \mathcal{C}'$ such that $\mathcal{G} \circ \mathcal{F}$ is isomorphic to $\mathbb{1}_{\mathcal{C}'}$ via a persistence natural transformation of shift 0, whose inverse also has shift 0, and the analogous condition holds for $\mathcal{F} \circ \mathcal{G}$ too. We will say that $\mathcal{C}'$ and $\mathcal{C}''$ are persistence equivalent or equivalent as persistence categories.

Standard arguments show that a persistence functor $\mathcal{F} : \mathcal{C}' \to \mathcal{C}''$ is an equivalence of persistence categories if and only if it is full and faithful (in the obvious persistence sense) and for every object $Y \in \mathrm{Obj}(\mathcal{C}'')$ there exists $X \in \mathrm{Obj}(\mathcal{C}')$ such that Y is 0-isomorphic to $\mathcal{F}(X)$ (i.e., the latter two objects are isomorphic in the 0-level subcategory $\mathcal{C}''_0$ of $\mathcal{C}''$).

2.2.3 Shift functors

The role of shift functors, to be introduced below, is to allow morphisms of arbitrary shift (as well as r-equivalences) to be represented as morphisms of shift 0, at the cost

of "shifting" the domain (or the target); see Remark 2.19. This turns out to be very helpful in the study of triangulation for persistence categories.

View the real axis $\mathbb{R}$ as a strict monoidal category $(\mathbb{R}, +)$ induced by the additive group structure of $\mathbb{R}$. In other words, $\mathrm{Obj}(\mathbb{R}) = \{x \mid x \in \mathbb{R}\}$ and, for any $x, y \in \mathrm{Obj}(\mathbb{R})$, $\hom_{\mathbb{R}}(x, y) = \{\eta_{x,y}\}$, such that $\eta_{x,x} = \mathbb{1}_x$ and, for any $x, y, z \in \mathbb{R}$, $\eta_{y,z} \circ \eta_{x,y} = \eta_{x,z}$. In particular, $\eta_{x,y} \circ \eta_{y,x} = \mathbb{1}_y$ and $\eta_{y,x} \circ \eta_{x,y} = \mathbb{1}_x$; hence each morphism $\eta_{x,y}$ is an isomorphism. The monoidal structure is defined by $\oplus(x, y) := x + y$ on objects and, for any two morphisms $(\eta_{r,r'}, \eta_{s,s'})$, we have $\oplus(\eta_{r,r'}, \eta_{s,s'}) := \eta_{r+s,r'+s'}$.

Definition 2.15. Let $\mathcal{C}$ be a persistence category. A *shift functor* on $\mathcal{C}$ is a strict monoidal functor $\Sigma : (\mathbb{R}, +) \to \mathcal{P}\mathrm{End}(\mathcal{C})$ such that $\Sigma(\eta_{x,y}) : \Sigma(x) \to \Sigma(y)$ is a natural transformation of shift $\lceil \Sigma(\eta_{x,y}) \rceil = y - x$ for any $x, y \in \mathbb{R}$ and $\eta_{x,y} \in \hom_{\mathbb{R}}(x, y)$.

For later use, we write $\Sigma^r := \Sigma(r) \in \mathcal{P}\mathrm{End}(\mathcal{C})$ and, for brevity, we denote $\Sigma(\eta_{r,s})$ by $\eta_{r,s}$ for $r, s \in \mathbb{R}$. We let $(\eta_{r,s})_A$ be the respective morphism $\Sigma^r A \to \Sigma^s A$.

Remark 2.16. Since Σ is a strict monoidal functor, it preserves the monoidal product. Therefore, $\Sigma^s \circ \Sigma^r = \Sigma^{r+s}$ and $\Sigma^0 = \mathbb{1}$. Moreover, since each $\eta_{r,s}$ is an isomorphism in $(\mathbb{R}, +)$, the corresponding natural transformation $\eta_{r,s}$ is a natural isomorphism. We also have $\Sigma^r(\eta_{s,s'})_A = (\eta_{s+r,s'+r})_A$ for each object A in $\mathcal{C}$ and all $r, s, s' \in \mathbb{R}$.

In particular, this implies that for any $Y, A \in \mathrm{Obj}(\mathcal{C})$ and $\alpha \in \mathbb{R}$, we have an isomorphism

$$\hom^{\alpha}_{\mathcal{C}}(Y, A) \xrightarrow{(\eta_{0,r})_A \circ} \hom^{\alpha+r}_{\mathcal{C}}(Y, \Sigma^r A). \tag{2.8}$$

Similarly, for any $A, X \in \mathrm{Obj}(\mathcal{C})$ and $\alpha \in \mathbb{R}$, we have an isomorphism

$$\hom^{\alpha-r}_{\mathcal{C}}(\Sigma^r A, X) \xrightarrow{\circ(\eta_{0,r})_A} \hom^{\alpha}_{\mathcal{C}}(A, X). \tag{2.9}$$

Further, for any $A, B \in \mathrm{Obj}(\mathcal{C})$, the isomorphisms (2.8) and (2.9) imply the existence of an isomorphism

$$\hom^{\alpha+s-r}_{\mathcal{C}}(\Sigma^r A, \Sigma^s B) \simeq \hom^{\alpha}_{\mathcal{C}}(A, B). \tag{2.10}$$

In particular, when $r = s$, we get a canonical isomorphism

$$\Sigma^r : \hom^{\alpha}_{\mathcal{C}}(A, B) \to \hom^{\alpha}_{\mathcal{C}}(\Sigma^r A, \Sigma^r B). \tag{2.11}$$

Finally, the diagrams (2.6) and (2.7) imply that the following diagrams, obtained by setting $\mathcal{F} = \Sigma^0$ and $\mathcal{G} = \Sigma^r$, are commutative for any $\alpha \leq \beta$:

$$\begin{array}{ccc} \hom_{\mathcal{C}}^{\alpha}(X, A) & \xrightarrow{(\eta_{0,r})_A \circ} & \hom_{\mathcal{C}}^{\alpha+r}(X, \Sigma^r A) \\ \downarrow{\scriptstyle i_{\alpha,\beta}} & & \downarrow{\scriptstyle i_{\alpha+r,\beta+r}} \\ \hom_{\mathcal{C}}^{\beta}(X, A) & \xrightarrow{(\eta_{0,r})_A \circ} & \mathrm{Mor}_{\mathcal{C}}^{\beta+r}(X, \Sigma^r A) \end{array} \tag{2.12}$$

and

$$\begin{array}{ccc} \hom_{\mathcal{C}}^{\alpha}(\Sigma^r A, X) & \xrightarrow{\circ(\eta_{0,r})_A} & \hom_{\mathcal{C}}^{\alpha+r}(A, X) \\ \downarrow{\scriptstyle i_{\alpha,\beta}} & & \downarrow{\scriptstyle i_{\alpha+r,\beta+r}} \\ \hom_{\mathcal{C}}^{\beta}(\Sigma^r A, X) & \xrightarrow{\circ(\eta_{0,r})_A} & \hom_{\mathcal{C}}^{\beta+r}(A, X) \end{array} \tag{2.13}$$

All the horizontal morphisms in (2.12) and (2.13) are isomorphisms, but the vertical morphisms (which are the persistence structure morphisms) are not necessarily so.

Assume that $\mathcal{C}$ is a persistence category (with persistence structure morphisms denoted by $i_{r,s}$) endowed with a shift functor Σ. To ease notation, for $A \in \mathrm{Obj}(\mathcal{C})$, $r \geq 0$, we consider $(\eta_{r,0})_A \in \hom_{\mathcal{C}}^{-r}(\Sigma^r A, A)$ and $(\eta_{0,-r})_A \in \hom_{\mathcal{C}}^{-r}(A, \Sigma^{-r} A)$, and we will denote below by η_r^A the maps

$$\eta_r^A = i_{-r,0}((\eta_{r,0})_A) \quad \text{or} \quad \eta_r^A = i_{-r,0}((\eta_{0,-r})_A). \tag{2.14}$$

Thus $\eta_r^A \in \hom_{\mathcal{C}}^0(\Sigma^r A, A)$ or $\eta_r^A \in \hom_{\mathcal{C}}^0(A, \Sigma^{-r} A)$, depending on the context. Note that there is no ambiguity of the notation η_r^A due to the canonical identification via Σ^r in (2.11). The notions discussed before, r-acyclicity, r-equivalence, and so forth, can be reformulated in terms of compositions with appropriate shift morphisms η_r^A.

The next lemma is a characterization of r-equivalence that follows easily from the diagrams (2.12) and (2.13).

Lemma 2.17. *Suppose that $f \in \hom^{\alpha}(A, B)$. Then $i_{\alpha,\alpha+r}(f) = 0$ for some $r \geq 0$ if and only if $f \circ \eta_r^A = 0$ in $\hom^{\alpha}(\Sigma^r A, B)$ and (equivalently) if and only if $\eta_r^B \circ f = 0$ in $\hom^{\alpha}(A, \Sigma^{-r} B)$.*

In particular, we easily see that for two morphisms $f, g \in \hom^{\alpha}(A, B)$, $f \simeq_r g$ if and only if $f \circ \eta_r^A = g \circ \eta_r^A$. Moreover, r-equivalence is preserved under shifts. Further, it is immediate to check that $f \in \hom^{\alpha}(A, B)$ and $g \in \hom^{\beta}(A, B)$ are ∞-equivalent if and only if there exist $r, r' \geq 0$ with $\alpha + r = \beta + r'$ such that

$$f \circ \eta_r^A = g \circ \eta_{r'}^A \quad \text{in } \hom^{\alpha+r}(A, B),$$

where we identify both $\hom^{\alpha}(\Sigma^r A, B)$ and $\hom^{\beta}(\Sigma^{r'} A, B)$ with $\hom^{\alpha+r}(A, B)$ through the canonical isomorphisms in Remark 2.16.

Here is a similar characterization of r-acyclicity.

Lemma 2.18. *$K \simeq_r 0$ is equivalent to each of the following:*

(i) $\eta_r^K = 0$.

(ii) $i_{\alpha,\alpha+r} : \hom^\alpha(A, K) \to \hom^{\alpha+r}(A, K)$ *vanishes for any $\alpha \in \mathbb{R}$ and A.*

(iii) $i_{\alpha,\alpha+r} : \hom^\alpha(K, A) \to \hom^{\alpha+r}(K, A)$ *vanishes for any $\alpha \in \mathbb{R}$ and A.*

Proof. Point (i) is an immediate consequence of the definition of r-acylics in Definition 2.8 and of Lemma 2.17 applied for $A, B = K$, $f = \mathbb{1}_K$. We now prove (ii). The proof of (iii) is similar and will be omitted. It is obvious that (ii) implies $K \simeq_r 0$ by specializing to $A = K, \alpha = 0$ and applying $i_{0,r}$ to $\mathbb{1}_K$. To prove the converse, we first use the diagram (2.12) to deduce that the map $i_{\alpha,\alpha+r}$ factors as

$$\hom^\alpha(A, K) \xrightarrow{(i_{-r,0}(\eta_{0,-r})_K)\circ} \hom^\alpha(A, \Sigma^{-r}K) \xrightarrow{(\eta_{-r,0})_K\circ} \hom^{\alpha+r}(A, K)\ .$$
$$\qquad\qquad\qquad\qquad\qquad\qquad i_{\alpha,\alpha+r}$$

Therefore, since $(\eta_{-r,0})_K\circ$ is an isomorphism, for any $f \in \hom^\alpha(A, K)$ we have $i_{\alpha,\alpha+r}(f) = 0$ if and only if $i_{-r,0}(\eta_{0,-r})_K \circ f = 0$. From point (i) we know that this relation is true for $f = \mathbb{1}_K$. Now, for any $f \in \hom^\alpha(A, K)$, we write $f = \mathbb{1}_K \circ f$ and conclude that $i_{-r,0}(\eta_{0,-r})_K \circ f = i_{-r,0}(\eta_{0,-r} \circ \mathbb{1}_K) \circ f = 0$. ∎

In particular, we see that K is r-acyclic if and only if any of its shifts $\Sigma^\alpha K$ is so.

Remark 2.19. (a) Assume that $\mathcal{C}$ is a persistence category endowed with a shift functor Σ and that $f_1, f_2 \in \hom^\alpha_{\mathcal{C}}(A, B)$. Then, for all practical purposes, we may replace f_i with the $\mathcal{C}_0$ morphisms $\widetilde{f_i} \in \hom^0_{\mathcal{C}}(\Sigma^\alpha A, B)$, where $\widetilde{f_i} = f_i \circ (\eta_{\alpha,0})_A$. The property $f_1 \simeq_r f_2$ is equivalent to

$$\widetilde{f_1} \circ \eta_r^{\Sigma^\alpha A} \simeq_0 \widetilde{f_2} \circ \eta_r^{\Sigma^\alpha A},$$

which is a relation in $\mathcal{C}_0$.

(b) Shift functors are natural in many geometric examples. Nonetheless, for a given persistence category $\mathcal{C}$, the existence of a shift functor Σ on $\mathcal{C}$ is a constraining additional structure. In particular, a persistence category $\mathcal{C}$ endowed with a shift functor Σ contains considerable redundant information. Indeed, the isomorphism (2.10) implies that all the morphisms in $\mathcal{C}$ are determined by the morphisms in $\mathcal{C}_0$ together with Σ. In other words, given a category $\mathcal{C}_0$ endowed with a shift functor Σ (appropriately defined), one can define a persistence category $\mathcal{C}$ with the same objects as $\mathcal{C}_0$ by using (2.10) to define morphisms of arbitrary shifts out of the morphisms in $\mathcal{C}_0$. We will see such an example in Section 2.2.4.

(c) There is an obvious way to formally complete any persistence category $\mathcal{C}$ to a larger persistence category $\widetilde{\mathcal{C}}$ that is endowed with a canonical shift functor.

This is achieved by formally adding objects $\Sigma^r X$ for each $r \in \mathbb{R}$ and $X \in \mathrm{Obj}(\mathcal{C})$, and defining morphisms such that the relations in Remark 2.16 are satisfied. In view of this and of the redundancy at point (b), one could prefer to replace the notion of a persistence category with a structure consisting of a category – corresponding to $\mathcal{C}_0$ – and a shift functor. This leads to an equivalent formalism. We stick in this memoir with the formalism of persistence categories as introduced in Definition 2.6, as we found it easiest to handle in algebraic manipulations and because it corresponds naturally to most of our geometric examples.

2.2.4 An example of a persistence category

We give an example of a persistence category that is constructed from persistence $\mathbf{k}$-modules. To some extent, this is the motivation of the definition of a persistence category. Recall that for a persistence $\mathbf{k}$-module $\mathbb{V}$, $\mathbb{V}[r]$ denotes another persistence $\mathbf{k}$-module which comes from an r-shift from $\mathbb{V}$ in the sense that

$$\mathbb{V}[r]_t = \mathbb{V}_{r+t} \quad \text{and} \quad \iota^{\mathbb{V}[r]}_{s,t} = \iota^{\mathbb{V}}_{s+r,t+r}.$$

A persistence morphism $f : \mathbb{V} \to \mathbb{W}$ is an $\mathbb{R}$-family of morphisms $f = \{f_t\}$ that commutes with the persistence structure maps of $\mathbb{V}$ and $\mathbb{W}$, i.e., $f_t \circ \iota^{\mathbb{V}}_{s,t} = \iota^{\mathbb{W}}_{s,t} \circ f_s$. Similarly, one can define an r-shifted persistence morphism $f[r]$, where $(f[r])_t = f_{r+t}$.

Let $\mathcal{P}\mathrm{Mod}_{\mathbf{k}}$ be the category of persistence $\mathbf{k}$-modules; then we claim that $\mathcal{P}\mathrm{Mod}_{\mathbf{k}}$ can be enriched to be a persistence category $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$. Indeed, let $\mathrm{Obj}(\mathcal{C}^{\mathcal{P}}\mathrm{Mod}_{\mathbf{k}}) = \mathrm{Obj}(\mathcal{P}\mathrm{Mod}_{\mathbf{k}})$, and for objects $\mathbb{V}, \mathbb{W}$ in $\mathcal{P}\mathrm{Mod}_{\mathbf{k}}$, define

$$\hom_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\mathbb{V}, \mathbb{W}) := \left\{\{\hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}, \mathbb{W}[r])\}_{r\in\mathbb{R}}; \{i_{r,s}\}_{r,s\in\mathbb{R}, r\le s}\right\}. \tag{2.15}$$

Here, $\hom^r_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\mathbb{V}, \mathbb{W}) = \hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}, \mathbb{W}[r])$, and $\hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\cdot,\cdot)$ consists of persistence morphisms. For any $r \le s$, the well-defined persistence morphism $\iota^{\mathbb{W}}_{\cdot+r,\cdot+s} : \mathbb{W}[r] \to \mathbb{W}[s]$ induces structure maps $i_{r,s} := \iota^{\mathbb{W}}_{\cdot+r,\cdot+s}\circ$ in (2.15). Moreover, the composition $\circ_{(r,s)} : \hom^r_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\mathbb{U}, \mathbb{V}) \times \hom^s_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\mathbb{V}, \mathbb{W}) \to \hom^{r+s}_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\mathbb{U}, \mathbb{W})$ is defined by

$$(f, g) \mapsto g[r] \circ f,$$

where we use the identification $\hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}, \mathbb{W})[r] = \hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}[r], \mathbb{W}[r])$ for any $r \in \mathbb{R}$. Moreover, for the following diagram, where $r \le r'$ and $s \le s'$:

$$\begin{array}{ccc}
\hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{U}, \mathbb{V}[r]) \times \hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}, \mathbb{W}[s]) & \xrightarrow{\circ_{(r,s)}} & \hom_{\mathcal{P}}(\mathbb{U}, \mathbb{W}[r+s]) \\
\downarrow{\scriptstyle i_{r,r'}\times i_{s,s'}} & & \downarrow{\scriptstyle i_{r+s,r'+s'}} \\
\hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{U}, \mathbb{V}[r']) \times \mathrm{Hom}_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}, \mathbb{W}[s']) & \xrightarrow{\circ_{(r',s')}} & \hom_{\mathcal{P}}(\mathbb{U}, \mathbb{W}[r'+s'])
\end{array} \tag{2.16}$$

we have

$$\begin{aligned}(\circ_{(r',s')} \circ (i_{r,r'} \times i_{s,s'}))(f,g) &= \circ_{(r',s')}(i_{r,r'}(f), i_{s,s'}(g)) \\ &= i_{s,s'}(g)[r'] \circ i_{r,r'}(f) \\ &= (\iota^{\mathbb{W}}_{\cdot+s,\cdot+s'} \circ g)[r'] \circ (\iota^{\mathbb{W}}_{\cdot+r,\cdot+r'} \circ f) \\ &= \iota^{\mathbb{W}}_{\cdot+r'+s,\cdot+r'+s'} \circ g[r'] \circ \iota^{\mathbb{V}}_{\cdot+r,\cdot+r'} \circ f \\ &= \iota^{\mathbb{W}}_{\cdot+r'+s,\cdot+r'+s'} \circ \iota^{\mathbb{W}}_{\cdot+r+s,\cdot+r'+s} \circ g[r] \circ f \\ &= \iota^{\mathbb{W}}_{\cdot+r+s,\cdot+r'+s'} \circ (g[r] \circ f) \\ &= (\iota_{r+s,r'+s'} \circ \circ_{(s,t)})(f,g),\end{aligned}$$

where the fifth equality is due to the fact that g is a persistence morphism (so, in particular, it commutes with the persistence structure maps). Therefore, the diagram (2.16) is commutative and $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$ is a persistence category in the sense of Definition 2.6.

Since $\hom^0_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\mathbb{V},\mathbb{W}) = \hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V},\mathbb{W})$, we have $\mathcal{C}_0^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}} = \mathcal{P}\mathrm{Mod}_{\mathbf{k}}$. An example of a persistence endofunctor on $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$, denoted by $\Sigma^\alpha : \mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}} \to \mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$, is defined by

$$\Sigma^\alpha(\mathbb{V}) := \mathbb{V}[-\alpha] \quad \text{and} \quad \Sigma^\alpha(f) = f[-\alpha] \tag{2.17}$$

for any $\alpha \in \mathbb{R}$. It is immediate to see that Σ^α is a persistence endofunctor on $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$ for any $\alpha \in \mathbb{R}$ in the sense of Definition 2.11.

We now define a shift functor on $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$, denoted by

$$\Sigma : (\mathbb{R},+) \to \mathcal{P}\mathrm{Fun}(\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}),$$

by

$$\Sigma(\alpha) := \Sigma^\alpha \text{ as defined in (2.17)} \quad \text{and} \quad \eta_{\alpha,\beta} = \mathbb{1}_{\cdot[-\alpha]}$$

for any $\alpha, \beta \in \mathbb{R}$. Indeed, evaluate $\eta_{\alpha,\beta}$ on any object $\mathbb{V}$:

$$\begin{aligned}(\eta_{\alpha,\beta})_{\mathbb{V}} = \mathbb{1}_{\mathbb{V}[-\alpha]} &\in \hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}[-\alpha], \mathbb{V}[-\alpha]) \\ &= \hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}[-\alpha], \mathbb{V}[-\beta+\beta-\alpha]) \\ &= \hom^{\beta-\alpha}_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\mathbb{V}[-\alpha], \mathbb{V}[-\beta]) \\ &= \hom^{\beta-\alpha}_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\Sigma^\alpha\mathbb{V}, \Sigma^\beta\mathbb{V}).\end{aligned}$$

In other words, $\eta_{\alpha,\beta}$ is a persistence natural transformation of shift $\beta-\alpha$ as in Definition 2.12. Therefore, Σ defines a shift functor on $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$.

Finally, for each $r \geq 0$, recall that the notation $\eta^{\mathbb{V}}_r$ in (2.14) denotes the composition $i_{-r,0} \circ (\eta_{r,0})_{\mathbb{V}}$. In particular, $\eta^{\mathbb{V}}_r \in \hom^0_{\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}}(\Sigma^r\mathbb{V},\mathbb{V}) = \hom_{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}(\mathbb{V}[-r],\mathbb{V})$ equals the composition

$$\mathbb{V}[-r] \xrightarrow{\mathbb{1}_{\mathbb{V}[-r]}} \mathbb{V}[-r] \xrightarrow{\iota^{\mathbb{V}}_{\cdot+r,\cdot}\circ} \mathbb{V},$$

which is just $\iota^{\mathbb{V}}_{-r,\cdot}$, the persistence structure maps of $\mathbb{V}$. Assume that objects in $\mathcal{P}\mathrm{Mod}_{\mathbf{k}}$ admit sufficient regularities so that they can be equivalently described via barcodes (see [25]). In this case, by Lemma 2.18 (i), the r-acyclic objects in $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$ are precisely those persistence $\mathbf{k}$-modules with only bars of length at most r in their barcodes (see [57–59]).

Remark 2.20. (a) The way that the category $\mathcal{P}\mathrm{Mod}_{\mathbf{k}}$ is enriched to $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$ above is also investigated in the recent work [13, Section 10]. In particular, the morphism set defined in (2.15) coincides with the enriched morphism set in [13, Proposition 10.2], and $\mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbf{k}}}$ is similar to $\underline{\mathrm{Mod}}^P_R$ in [13, Proposition 10.3] (when taking $R = \mathbf{k}$ and $P = \mathbb{R}$).

(b) The notion of a persistence category endowed with shift functors is very natural in persistence considerations, as already mentioned in point (a) above. The same notion appears in [51] under the name of a *locally persistent category*, and the 0-level (from Definition 2.9) appears in that work, where it is called the *underlying category* of the respective locally persistent category. The definition of interleaving of persistence modules adapts trivially to this context – of persistence categories or, equivalently, locally persistent categories – to provide a (pseudo)distance, possibly infinite, on the objects of such a category, as in Definition 2.84. The work [51] analyzes and establishes key properties – for instance, completeness – for interleaving distances of this sort under certain assumptions – such as existence of products, or co-products, or a model structure, or existence of limits – on the 0-level category.

Starting from Section 2.3, we focus on properties of persistence categories with 0-levels that have the structure of triangulated categories. In practice, this means that they are often homotopy categories of other categories. From this perspective, while we work, in some sense, at the homotopy level, [51] is geared towards considering 0-levels that directly have a model category structure. In our case, the triangulation is tied to the persistence structure by some simple axioms. If $\mathcal{C}$ is such a category, called a *triangulated persistence category (TPC)*, then we will see that the ∞-level is endowed with a categorical weight induced by the persistence structure, and the general machinery in Section 2.1 leads to a class of fragmentation (pseudo)metrics on the objects of $\mathcal{C}$. These fragmentation metrics extend, on the one hand, interleaving-type metrics and, on the other hand, complexity measurements such as those mentioned in Remark 1.2; they are also similar to classical notions in topology such as cone-length [19]. In Section 3 we show that certain derived Fukaya categories admit TPC refinements and thus their objects are endowed with persistence fragmentation metrics. The interest of this class of fragmentation metrics in this symplectic context is that, under favourable geometric assumptions, these metrics are both non-degenerate and finite while interleaving-type distances take infinite values.

Some elementary relations between interleaving and the rest of the algebraic machinery in the memoir appear in Section 2.4.3.4. Moreover, it is likely that, in some

cases, some of the deeper properties of the interleaving distances discussed in [51] can be related in more substantial ways to our fragmentation metrics; however, we will not pursue these questions here.

2.3 Triangulated persistence categories

This section is central for the rest of the memoir. It investigates triangulation properties in the context of persistence categories. We start with two key definitions in Section 2.3.1: Definition 2.21, which introduces the notion of a triangulated persistence category (TPC) – a persistence category $\mathcal{C}$ with a shift functor whose 0-level $\mathcal{C}_0$ is triangulated – and Definition 2.26, which introduces the notion of an r-isomorphism. We then discuss a number of useful properties of r-isomorphisms. These properties are, in some sense, "shift"-controlled analogs of properties that appear when defining the Verdier localization of a triangulated category. Indeed, in Section 2.3.2 we see that the acyclics of finite order in $\mathcal{C}$ form a triangulated subcategory $\mathcal{AC}_0$ of $\mathcal{C}_0$, and that the Verdier localization of $\mathcal{C}_0$ with respect to this subcategory is the ∞-level category $\mathcal{C}_\infty$ of $\mathcal{C}$, which is therefore itself triangulated. The main aim of our algebraic formalism is to construct a notion of weighted exact triangles in $\mathcal{C}$ – and this is pursued in Section 2.3.3, in particular in Definition 2.3.3. We then discuss in Section 2.3.4 associated fragmentation pseudometrics.

2.3.1 Main definitions

We will consistently use below the characterization of r-equivalence in Lemma 2.17 as well as that of r-acyclics in Lemma 2.18.

Definition 2.21. A *triangulated persistence category* is a persistence category $\mathcal{C}$ endowed with a shift functor Σ such that the following three conditions are satisfied:

(i) The 0-level category $\mathcal{C}_0$ is triangulated with a translation automorphism denoted by T. Note that, in particular, $\mathcal{C}_0$ is additive, and we further assume that the restriction of the persistence structure of $\mathcal{C}$ to $\mathcal{C}_0$ is compatible with the additive structure on $\mathcal{C}_0$ in the obvious way. Specifically, this means that $\hom^r_{\mathcal{C}}(A \oplus B, C) = \hom^r_{\mathcal{C}}(A, C) \oplus \hom^r_{\mathcal{C}}(B, C)$ for all $r \leq 0$, and the persistence maps $i_{r,s}$, $r \leq s \leq 0$, are compatible with this splitting. The same also holds for $\hom^r_{\mathcal{C}}(A, B \oplus C)$.

(ii) The restriction of Σ^r to $\mathrm{End}(\mathcal{C}_0)$ is a triangulated endofunctor of $\mathcal{C}_0$ for each $r \in \mathbb{R}$. Note that each of the functors Σ^r, being a triangulated functor, is also assumed to be additive. We further assume that all the natural transformations $\eta_{r,s} : \Sigma^r \to \Sigma^s$, $s, r \in \mathbb{R}$, are compatible with the additive structure on $\mathcal{C}_0$.

(iii) For any $r \geq 0$ and any $A \in \mathrm{Obj}(\mathcal{C})$, the morphism $\eta_r^A : \Sigma^r A \to A$ defined in (2.14) embeds into an exact triangle of $\mathcal{C}_0$,

$$\Sigma^r A \xrightarrow{\eta_r^A} A \to K \to T\Sigma^r A,$$

such that K is r-acyclic.

Example 2.22. The fundamental example of a triangulated persistence category is provided by the homotopy category of filtered (co)-chain complexes over a field $\mathbf{k}$, $H^0\mathcal{FK}_{\mathbf{k}}$. The objects are filtered cochain complexes (C, ∂) over $\mathbf{k}$, with

$$C : \cdots \subset C^{\leq\alpha} \subset C^{\leq\beta} \subset \cdots \quad (\alpha \leq \beta \in \mathbb{R}),$$

and ∂ does not increase filtration; hence each $C^{\leq\alpha}$ is itself a cochain complex – a more complete description is given in Section 2.5.2. The morphisms are homotopy classes of filtered chain maps

$$\hom^r(C, C') = \{f : C \to C' \mid f \text{ is a chain map, } f(C^{\leq\alpha}) \subset (C')^{\leq\alpha+r}\}/\simeq_r,$$

where the relation $\simeq_r$ is cochain homotopy via a homotopy $h : C^* \to C^{*-1}$ such that $h(C^{\leq\alpha}) \subset (C')^{\leq\alpha+r}$. The translation functor is defined as usual by translating degree (and keeping the filtration unchanged), namely $(TC)^i = C^{i+1}$, and with the obvious action on morphisms. The shift functor acts on objects by $[\Sigma^r C]^{\leq\alpha} = C^{\leq\alpha-r}$ with the obvious differential and the obvious action on morphisms. The 0-level of $H^0\mathcal{FK}_{\mathbf{k}}$, $[H^0\mathcal{FK}_{\mathbf{k}}]_0$, is the subcategory with the same objects but whose morphisms come only from filtration-preserving chain maps. This is a triangulated category because, for chain-preserving maps, the mapping-cone construction is filtration preserving. The r-acyclics in this case are filtered complexes C such that $\mathbb{1}_C$ is chain homotopic to 0 through a chain homotopy that shifts filtration by at most r.

Note that we also have the (full) subcategory $\mathcal{FK}_k^{\mathrm{fg}} \subset \mathcal{FK}_{\mathbf{k}}$ of *finitely generated* filtered cochain complexes, which is also a TPC. The category $[H^0\mathcal{FK}_{\mathbf{k}}^{\mathrm{fg}}]_\infty$ is equivalent to the usual homotopy category of finitely generated cochain complexes.

Remark 2.23. (a) Given that $\mathcal{C}_0$ is triangulated, the functors $\hom^s_{\mathcal{C}}(X, -)$ and $\hom^s_{\mathcal{C}}(-, X)$ are exact for $s = 0$. This property, together with the fact that Σ^s is a triangulated functor for all s and the relations in Remark 2.16, implies that these functors are exact *for all* $s \in \mathbb{R}$.

(b) Condition (ii) requires in particular that Σ and T commute. Thus, $T\Sigma^r X = \Sigma^r TX$ for each object X, and for any $f \in \hom^0_{\mathcal{C}}(A, B)$ we have $\Sigma^r Tf = T\Sigma^r f$. Additionally, each Σ^r preserves the additive structure of $\mathcal{C}_0$, and it takes each exact triangle in $\mathcal{C}_0$ to an exact triangle. Moreover, the assumptions above imply that we have canonical isomorphisms

$$\hom^r_{\mathcal{C}}(A \oplus B, C) \cong \hom^r_{\mathcal{C}}(A, C) \oplus \hom^r_{\mathcal{C}}(B, C) \quad \forall r \in \mathbb{R},$$

and the persistence maps $i_{r,s}$ are compatible with these isomorphisms. The same holds also for $\hom^r_{\mathcal{C}}(A, B \oplus C)$. Finally, the maps η^A_r, $A \in \mathrm{Obj}(\mathcal{C}_0)$, $r \geq 0$, are compatible with the additive structure on $\mathcal{C}_0$.

Notice also that the functor T extends from $\mathcal{C}_0$ to a functor on $\mathcal{C}$. Indeed, T is already defined on all the objects of $\mathcal{C}$ as well as on all the morphisms of shift 0. For $f \in \hom^r(A, B)$, we define $Tf = T(f \circ (\eta_{r,0})_A) \circ (\eta_{0,r})_{TA}$. It is easily seen that with this definition T is indeed a functor and it immediately follows that $T((\eta_{r,s})_A) = (\eta_{r,s})_{TA}$ for all objects A in $\mathcal{C}$ and $r, s \in \mathbb{R}$. Further, by using the identifications in Remark 2.16, it also follows that T is a persistence functor. In particular, we have $\eta^{TA}_r = T\eta^A_r$ for each object A in $\mathcal{C}$.

(c) Given that 0-isomorphisms preserve r-acyclicity, as noted in Section 2.2.1, condition (iii) in Definition 2.21 does not depend on the specific extension of η^A_r to an exact triangle.

(d) In a way similar to Remark 2.19 (b), the data encoded in a triangulated persistence category is determined by the triangulated category $\mathcal{C}_0$ together with an appropriate shift functor $\Sigma : (\mathbb{R}, +) \to \mathrm{End}(\mathcal{C}_0)$. From this data it is easy to define a triangulated persistence category $\mathcal{C}$ with the same objects as $\mathcal{C}_0$, that has $\mathcal{C}_0$ as its 0-level and with morphisms endowed with a persistence structure such that (2.12) and (2.13) are satisfied with respect to the given shift functor Σ. We do not give further details here, but we will see such an example in Section 2.5.4.

It is clear that TPCs form a category with respect to persistence functors that respect the additional structure. The appropriate notion is formalized below.

Definition 2.24. Let $\mathcal{C}'$ and $\mathcal{C}''$ be two TPCs. A persistence functor $\mathcal{F} : \mathcal{C}' \to \mathcal{C}''$ is called a *TPC functor* if it satisfies the following conditions:

(i) $\mathcal{F}$ is compatible with the shift functors $\Sigma_{\mathcal{C}'}$, $\Sigma_{\mathcal{C}''}$ of the two categories, namely $\mathcal{F} \circ \Sigma^r_{\mathcal{C}'} = \Sigma^r_{\mathcal{C}''} \circ \mathcal{F}$ for all $r \in \mathbb{R}$, and $(\eta_{r,s})_{\mathcal{F}(A)} = \mathcal{F}((\eta_{r,s})_A)$ for all $A \in \mathrm{Obj}(\mathcal{C}')$ and all r, s.

(ii) The 0-level $\mathcal{F}|_{\mathcal{C}'_0} : \mathcal{C}'_0 \to \mathcal{C}''_0$ of the functor $\mathcal{F}$ is triangulated.

In the definition above, the fact that $\mathcal{F}|_{\mathcal{C}'_0}$ maps $\mathcal{C}'_0$ to $\mathcal{C}''_0$ follows from the assumption that $\mathcal{F}$ is a persistence functor. Modifying Definition 2.14, we now give the definition of an equivalence between TPCs.

Definition 2.25. Let $\mathcal{C}'$ and $\mathcal{C}''$ be two TPCs. A TPC functor $\mathcal{F} : \mathcal{C}' \to \mathcal{C}''$ is called a *TPC equivalence* if there exists a TPC functor $\mathcal{G} : \mathcal{C}'' \to \mathcal{C}'$ such that both $\mathcal{F} \circ \mathcal{G}$ and $\mathcal{G} \circ \mathcal{F}$ are isomorphic to the respective identity functors via persistence natural transformations of shift 0.

Standard results in triangulated categories (e.g., [39, Section 1.2]) imply that a TPC functor $\mathcal{F} : \mathcal{C}' \to \mathcal{C}''$ is a TPC equivalence if and only if it is an equivalence of persistence categories.

Definition 2.26. Let $\mathcal{C}$ be a triangulated persistence category. A map $f \in \hom^0_{\mathcal{C}}(A, B)$ is said to be an *r-isomorphism* (from A to B) if it embeds into an exact triangle in $\mathcal{C}_0$

$$A \xrightarrow{f} B \to K \to TA$$

such that $K \simeq_r 0$. In this case, we write $f : A \simeq_r B$.

Remark 2.27. (a) If f is an r-isomorphism, then f is an s-isomorphism for any $s \geq r$. It is not difficult to check, and we will see this explicitly in Remark 2.30, that for $r = 0$ this definition is equivalent to the notion of 0-isomorphism introduced before (namely, isomorphism in the category $\mathcal{C}_0$).

(b) The relation $T(\eta_r^K) = \eta_r^{TK}$ implies that TK is r-acyclic if and only if K is r-acyclic and, therefore, f is an r-isomorphism if and only if Tf is one.

(c) From Definition 2.21 (iii) we see that for any $r \geq 0$ and $A \in \mathrm{Obj}(\mathcal{C})$ we have

$$\eta_r^A : \Sigma^r A \simeq_r A.$$

Proposition 2.28. *Any triangulated persistence category $\mathcal{C}$ has the following properties:*

(i) *If $f : A \to B$ is an r-isomorphism, then there exist $\phi \in \hom^0_{\mathcal{C}}(B, \Sigma^{-r} A)$ and $\psi \in \hom^0_{\mathcal{C}}(\Sigma^r B, A)$ such that*

$$\phi \circ f = \eta_r^A \text{ in } \hom^0_{\mathcal{C}}(A, \Sigma^{-r} A) \quad \text{and} \quad f \circ \psi = \eta_r^B \text{ in } \hom^0_{\mathcal{C}}(\Sigma^r B, B).$$

The map ψ is called a right r-inverse *of f and ϕ is a* left r-inverse *of f. They satisfy $\Sigma^r \phi \simeq_r \psi$.*

(ii) *If f is an r-isomorphism, then any two left r-inverses ϕ, ϕ' of f are themselves r-equivalent, and the same conclusion holds for right r-inverses.*

(iii) *If $f : A \simeq_r B$ and $g : B \simeq_s C$, then $g \circ f : A \simeq_{r+s} C$.*

Proof. (i) We first construct ϕ. In $\mathcal{C}_0$, the morphism $f : A \to B$ embeds into an exact triangle $A \xrightarrow{f} B \xrightarrow{g} K \xrightarrow{h} TA$ with $K \simeq_r 0$. Using the fact that Σ and T commute, the following diagram is easily seen to be commutative:

$$\begin{array}{ccc} K & \xrightarrow{\quad h \quad} & TA \\ {\scriptstyle \eta_r^K}\downarrow & & \downarrow{\scriptstyle \eta_r^{TA}} \\ \Sigma^{-r} K & \xrightarrow{\Sigma^{-r} h} & \Sigma^{-r} TA \end{array}$$

Thus $\eta_r^{TA} \circ h = \Sigma^{-r} h \circ \eta_r^K = 0$, since K is r-acyclic (and so $\eta_r^K = 0$). By rotating exact triangles in $\mathcal{C}_0$ we obtain a new $\mathcal{C}_0$-exact triangle $K \xrightarrow{h} TA \xrightarrow{Tf} TB \xrightarrow{Tg} TK$

and consider the diagram below (in $\mathcal{C}_0$):

$$\begin{array}{ccccccc} K & \xrightarrow{h} & TA & \xrightarrow{Tf} & TB & \xrightarrow{Tg} & TK \\ \downarrow & & \downarrow{\scriptstyle \eta_r^{TA}} & & \downarrow{\scriptstyle \widetilde{\phi}} & & \downarrow \\ 0 & \longrightarrow & \Sigma^{-r}TA & \xrightarrow{\mathbb{1}} & \Sigma^{-r}TA & \longrightarrow & 0 \end{array}$$

The left square commutes, so we deduce the existence of a map

$$\widetilde{\phi} \in \hom^0_{\mathcal{C}}(TB, \Sigma^{-r}TA)$$

that makes the middle and right squares commutative. The desired left inverse of f is $\phi = T^{-1}\widetilde{\phi}$. A similar argument leads to the existence of ψ. We postpone the identity $\Sigma^r\phi \simeq_r \psi$ after the proof of (ii).

(ii) If ϕ, ϕ' are two left inverses of f then $(\phi - \phi') \circ f = 0$. Therefore

$$(\phi - \phi') \circ \eta_r^B = (\phi - \phi') \circ (f \circ \psi) = ((\phi - \phi') \circ f) \circ \psi = 0.$$

Lemma 2.17 implies that $\phi \simeq_r \phi'$. The same argument works for right inverses. We now return to the identity $\Sigma^r\phi \simeq_r \psi$ (using the notation from (i)). We have the following commutative diagram:

$$\Sigma^r B \xrightarrow{\psi} A \underset{\eta_r^B}{\overset{f}{\rightleftarrows}} B \xrightarrow{\phi} \Sigma^{-r}A\,, \qquad \eta_r^A : A \to B$$

Therefore, $\eta_r^A \circ \psi = \phi \circ \eta_r^B$. By the naturality properties of η we also have $\phi \circ \eta_r^B = \eta_r^A \circ \Sigma^r\phi$. Thus, by Lemma 2.17, $\Sigma^r\phi \simeq_r \psi$.

(iii) We will make use of the following lemma.

Lemma 2.29. *If $K \to K'' \to K' \to TK$ is an exact triangle in $\mathcal{C}_0$, $K \simeq_r 0$ and $K' \simeq_s 0$, then $K'' \simeq_{r+s} 0$.*

Proof of Lemma 2.29. We associate the following commutative diagram to the exact triangle in the statement:

$$\begin{array}{ccccc} \hom^\alpha(K'', K) & \longrightarrow & \hom^\alpha(K'', K'') & \longrightarrow & \hom^\alpha(K'', K') \\ \downarrow & & \downarrow{\scriptstyle v} & & \downarrow{\scriptstyle 0} \\ \hom^{\alpha+s}(K'', K) & \xrightarrow{k} & \hom^{\alpha+s}(K'', K'') & \xrightarrow{h} & \hom^{\alpha+s}(K'', K') \\ {\scriptstyle 0}\downarrow & & \downarrow{\scriptstyle t} & & \\ \hom^{\alpha+r+s}(K'', K) & \xrightarrow{n} & \hom^{\alpha+r+s}(K'', K'') & & \end{array}$$

Here, the vertical morphisms are the persistence structure maps. The rightmost vertical map and the lower-leftmost vertical map are both 0 due to our hypothesis together with Lemma 2.18. The functor $\hom^{\alpha}(K'', -)$ is exact which implies that $t \circ v = 0$ and, again by Lemma 2.18, we deduce $K'' \simeq_{r+s} 0$. ∎

Returning to the proof of Proposition 2.28, point (iii) now follows immediately by using the octahedral axiom to construct the following commutative diagram in $\mathcal{C}_0$:

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \longrightarrow & K \\ \downarrow & & \downarrow{\scriptstyle g} & & \downarrow \\ A & \xrightarrow{g\circ f} & C & \longrightarrow & K'' \\ \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & K' & \longrightarrow & K' \end{array}$$

with exact rows and columns, and applying Lemma 2.29 to the rightmost column. ∎

Remark 2.30. (a) Points (i) and (ii) in Proposition 2.28 imply that the notion of 0-isomorphism $f : A \to B$, as given in Definition 2.26 for $r = 0$, is equivalent to an isomorphism in $\mathcal{C}_0$. In particular, for $r = 0$, f admits a unique inverse in $\hom^0_{\mathcal{C}}(B, A)$.

(b) Point (iii) in Proposition 2.28 shows that being r-isomorphic (for a fixed r) cannot be expected to be an equivalence relation on $\mathrm{Obj}(\mathcal{C})$ (unless $r = 0$).

Here are several useful additional results and corollaries.

The first is a version of the five-lemma in the TPC context.

Proposition 2.31. *Consider the following commutative diagram in $\mathcal{C}_0$:*

$$\begin{array}{ccccccc} A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & TA \\ {\scriptstyle u}\downarrow & & {\scriptstyle v}\downarrow & & {\scriptstyle w}\downarrow & & \downarrow{\scriptstyle Tu} \\ A' & \longrightarrow & B' & \longrightarrow & C' & \longrightarrow & TA' \end{array} \tag{2.18}$$

such that the two rows are exact triangles. If u is an r-isomorphism and v is an s-isomorphism, then:

(i) *There exists an $(r + s)$-isomorphism w making the diagram commutative.*

(ii) *Any w making the diagram commutative is a $3(r + s)$-isomorphism.*

Proof. Part (i) of the proposition is an easy consequence of Lemma 2.29 and the octahedral axiom in $\mathcal{C}_0$.

To show part (ii) we will use the following notation. For an object $Z \in \mathrm{Obj}(\mathcal{C})$, we denote by H_Z the functor $\hom_{\mathcal{C}}(Z, -) : \mathcal{C} \to \mathcal{C}^{\mathcal{P}\mathrm{Mod}_{\mathbb{k}}}$ (the target of this functor is an obvious enrichment of the category of persistence modules; see Section 2.2.4).

Similarly, denote by H_Z^* the (contravariant) functor $\hom_{\mathcal{C}}(-, Z)$. Further, if U and V are persistence modules and $F : U \to V$ is a morphism in ${\mathcal{C}}^{\mathcal{P}\mathrm{Mod}_{\Bbbk}}$, we say that F is:

- r-epi if for all $y \in V$ there exists $x \in U$ such that $F(x) = i_r(y)$. (Here and in what follows, i_r stands for the persistence structural map on $\hom_C(A, B)$ whose restriction to $\hom^{\alpha}(A, B)$ is the map $i_{\alpha,\alpha+r} : \hom_C^{\alpha}(A, B) \to \hom_C^{\alpha+r}(A, B)$, $\alpha \in \mathbb{R}$.)
- r-mono if for all $x \in U$ with $F(x) = 0$ we have $i_r(x) = 0$.

The proof is based on properties of right and left inverses that are contained in the following statement.

Lemma 2.32. *Let $\phi \in \hom_{\mathcal{C}_0}(M, N)$.*

(i) *ϕ admits a right r-inverse if and only if $H_Z(\phi)$ is r-epi for all $Z \in \mathrm{Obj}(\mathcal{C})$. The existence of a right r-inverse implies that $H_Z^*(\phi)$ is r-mono for all Z.*

(ii) *ϕ admits a left r-inverse if and only if $H_Z^*(\phi)$ is r-epi for all $Z \in \mathrm{Obj}(\mathcal{C})$. If such a left r-inverse exists, then $H_Z(\phi)$ is r-mono for all Z.*

(iii) *Consider a morphism of exact triangles as in (2.18). Assume that for every $Z \in \mathrm{Obj}(\mathcal{C})$, $H_Z(u)$ is r-epi and s-mono and that $H_Z(v)$ is r'-epi and s'-mono. Then $H_Z(w)$ is $(r' + r + s')$-epi and $(s + s' + r)$-mono for all Z.*

(iv) *If ϕ admits a right r-inverse and a left s-inverse, then ϕ is a $(r + s)$-isomorphism.*

Proof of Lemma 2.32. We start with (i). Assume that ϕ admits a right r-inverse $\psi \in \hom_{\mathcal{C}_0}(\Sigma^r N, M)$. Let $h \in \hom_{\mathcal{C}}^{\alpha}(Z, N)$. We have $h \circ \eta_r^Z = \eta_r^N \circ \Sigma^r h = \phi \circ (\psi \circ \Sigma^r h) = H_Z(\phi)(\psi \circ \Sigma^r h)$, with $\psi \circ \Sigma^r h \in \hom_{\mathcal{C}}^{\alpha}(\Sigma^r Z, M) = \hom_{\mathcal{C}}^{r+\alpha}(Z, M)$. Therefore

$$i_{\alpha,\alpha+r}(h) = (h \circ \eta_r^Z) \circ (\eta_{0,r})_Z = \phi \circ \big(\psi \circ \Sigma^r h \circ (\eta_{0,r})_Z\big).$$

It follows that $H_Z(\phi)$ is r-epi. Conversely, assume that $H_N(\phi)$ is r-epi. Then η_r^N is in the image of $H_N(\phi)$, which means that ϕ admits a right r-inverse. To finish with (i), let $k \in \hom_{\mathcal{C}}^{\alpha}(N, Z)$ be such that $k \circ \phi = 0$. We write $0 = k \circ \phi = k \circ \phi \circ \psi = k \circ \eta_r^N = \eta_r^Z \circ \Sigma^r k$, which means $i_r(k) = 0 \in \hom_{\mathcal{C}}^{r+\alpha}(N, Z)$. Thus $H_Z^*(\phi)$ is r-mono.

Point (ii) is entirely similar to (i).

For point (iii), we first notice that, by assumption, the maps $H_Z(u)$ and $H_Z(v)$ satisfy the epi and mono conditions with constants that are the same for all objects Z in $\mathcal{C}$. It is immediate to see that $H_Z(Tu)$ is r-epi (respectively, s-mono) if and only if $H_{T^{-1}Z}(u)$ is r-epi (and, respectively, s-mono). This implies that $H_Z(T^l u)$ is r-epi and s-mono, and that $H_Z(T^l v)$ is r'-epi and s'-mono for all $l \in \mathbb{Z}$ (and all Z). We now apply the exact functor H_Z to the diagram (2.18), and we obtain two long exact sequences of persistence modules related by comparison morphisms. The desired

conclusion follows by direct diagram chasing, as in the proof of the classical five-lemma.

For point (iv), we use the triangulated structure of $\mathcal{C}_0$ to obtain an object K and the following commutative diagram in $\mathcal{C}_0$:

$$\begin{array}{ccccccc} M & \xrightarrow{id} & M & \longrightarrow & 0 & \longrightarrow & TM \\ {\scriptstyle id}\downarrow & & {\scriptstyle \phi}\downarrow & & {\scriptstyle p}\downarrow & & {\scriptstyle id}\downarrow \\ M & \xrightarrow{\phi} & N & \longrightarrow & K & \longrightarrow & TM \end{array} \tag{2.19}$$

whose rows are exact triangles. Given that ϕ admits a right r-inverse, we deduce from (i) that $H_Z(\phi)$ is r-epi for all objects Z in $\mathcal{C}$. The existence of a left s-inverse implies, by (ii), that $H_Z(\phi)$ is also s-mono for all Z. We now use (iii) to deduce that $H_K(p)$ is $(r+s)$-epi. This implies that $i_{r+s}(id_K)$ is in the image of $H_K(p)$. But p is the null map, so $i_{r+s}(id_K) = 0$. Hence $K \simeq_{r+s} 0$. It follows that ϕ is an $(r+s)$-isomorphism. ∎

We now return to the proof of the second point of Proposition 2.31 with the notation and the assumptions there. We denote by K any object that completes the map w to an exact triangle

$$C \xrightarrow{w} C' \to K \to TC.$$

An r-isomorphism admits both right and left r-inverses. Thus, points (i) and (ii) of the lemma show that $H_Z(u)$ is r-epi and r-mono and that $H_Z(v)$ is s-epi and s-mono for all objects Z in $\mathcal{C}$. Point (iii) of the lemma then implies that $H_Z(w)$ is $(2s+r)$-epi and $(2r+s)$-mono for all Z. We now consider a diagram just as (2.19) but with $M = C$, $N = C'$, $\phi = w$, and we use point (iii) of the lemma to deduce that the map $H_K(p)$ is $3(r+s)$-epi, which means that K is $3(r+s)$-acyclic. ∎

Corollary 2.33. *If $f : A \to B$ is an r-isomorphism, then any right inverse $\psi \in \hom^0(\Sigma^r B, A)$ (given by* (i) *in Proposition 2.28) is a $2r$-isomorphism. The same conclusion holds for any left inverse.*

Proof. By the octahedral axiom (in $\mathcal{C}_0$), we have the commutative diagram

$$\begin{array}{ccccc} \Sigma^r B & \longrightarrow & \Sigma^r B & \longrightarrow & 0 \\ {\scriptstyle \psi}\downarrow & & {\scriptstyle \eta_r^B}\downarrow & & \downarrow \\ A & \longrightarrow & B & \longrightarrow & K \\ \downarrow & & \downarrow & & \downarrow \\ K'' & \longrightarrow & K' & \longrightarrow & K \end{array}$$

where $K'' \to K' \to K \to TK''$ is exact. By (iii) in Definition 2.21, $K' \simeq_r 0$. Therefore, by Lemma 2.29, $K'' \simeq_{2r} 0$ and thus ψ is $2r$-isomorphism. A similar argument applies to the left inverse of f. ∎

Remark 2.34. The fact that left and right inverses of r-isomorphisms are only $2r$-isomorphisms has significant impact on the various algebraic properties of TPCs. However, this seems unavoidable. For example, it is easy to construct examples of r-isomorphisms in the (homotopy) category of filtered cochain complexes that admit a unique right inverse which cannot be better than a $2r$-isomorphism.

The next consequence is immediate but useful, so we state it separately.

Corollary 2.35. *If $f : A \to B$ is an r-isomorphism, then for any $u, u' \in \hom^0_{\mathcal{C}}(B, C)$ such that $u \circ f = u' \circ f$, we have $u \simeq_r u'$, i.e., u and u' are r-equivalent. Similarly, if $v, v' \in \hom^0_{\mathcal{C}}(D, A)$ and $f \circ v = f \circ v'$, then $v \simeq_r v'$.*

Corollary 2.36. *Assume that the following diagram in $\mathcal{C}_0$:*

$$\begin{array}{ccccccc} K & \longrightarrow & A & \xrightarrow{\phi} & A' & \longrightarrow & TK \\ {\scriptstyle \mathbb{1}_K}\downarrow & & {\scriptstyle f}\downarrow & & {\scriptstyle f'}\downarrow & & {\scriptstyle \mathbb{1}_{TK}}\downarrow \\ K & \longrightarrow & B & \xrightarrow{\psi} & B' & \longrightarrow & TK \end{array}$$

is commutative, that the two rows are exact, and that $K \simeq_r 0$. Then the induced morphism f' is unique up to r-equivalence.

Proof. Since $K \simeq_r 0$, by definition, ϕ is an r-isomorphism. For any two induced morphisms $f'_1, f'_2 \in \hom^0_{\mathcal{C}}(A', B')$, we have $f'_1 \circ \phi = f'_2 \circ \phi = \psi \circ f$ and the conclusion follows from Corollary 2.35. ∎

Corollary 2.37. *Let $\phi : A \to A'$ be an r-isomorphism. Then for any $f \in \hom^0_{\mathcal{C}}(A, B)$, there exists $f' \in \hom^0_{\mathcal{C}}(A', \Sigma^{-r} B)$ such that the following diagram commutes in $\mathcal{C}_0$:*

$$\begin{array}{ccc} A & \xrightarrow{\phi} & A' \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f'} \\ B & \xrightarrow{\eta_r^B} & \Sigma^{-r} B \end{array}$$

Proof. Since $\phi : A \to A'$ is an r-isomorphism, there exists a left r-inverse denoted by $\psi : A' \to \Sigma^{-r} A$ such that $\psi \circ \phi = \eta_r^A$. Set $f' := \Sigma^{-r} f \circ \psi \in \hom^0_{\mathcal{C}}(A', \Sigma^{-r} B)$. ∎

Similar direct arguments lead to the next consequence.

Corollary 2.38. *Consider the following commutative diagram in $\mathcal{C}_0$:*

$$\begin{array}{ccc} A & \xrightarrow{\phi} & A' \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f'} \\ B & \xrightarrow{\phi'} & B' \end{array}$$

where $f \in \hom^0_{\mathcal{C}}(A, B)$, $f' \in \hom^0_{\mathcal{C}}(A'.B')$, *and* ϕ, ϕ' *are* r*-isomorphisms. Let* ψ, ψ' *be any left inverses of* ϕ, ϕ'*, respectively. Then the diagram*

$$\begin{array}{ccc} A' & \xrightarrow{\psi} & \Sigma^{-r} A \\ {\scriptstyle f'}\downarrow & & \downarrow{\scriptstyle \Sigma^{-r} f} \\ B' & \xrightarrow{\psi'} & \Sigma^{-r} B \end{array}$$

is r*-commutative, in the sense that* $\Sigma^{-r} f \circ \psi \simeq_r \psi' \circ f'$*. A similar conclusion holds for right inverses.*

2.3.2 Relation to Verdier localization

Proposition 2.39. *Let* $\mathcal{C}$ *be a triangulated persistence category and let* $\mathcal{AC}$ *be the full subcategory of* $\mathcal{C}$ *with objects the* r*-acyclic objects of* $\mathcal{C}$ *(for all* $r \geq 0$*).*

(i) *The category* $\mathcal{AC}$ *is a triangulated persistence category on its own, with 0-level denoted by* $\mathcal{AC}_0$ $(= (\mathcal{AC})_0)$*, the full subcategory of* $\mathcal{C}_0$ *whose objects are those of* $\mathcal{AC}$*.*

(ii) *The infinity level* $\mathcal{C}_\infty$ *of* $\mathcal{C}$ *coincides with the Verdier quotient (a.k.a. localization)* $\mathcal{C}_0/\mathcal{AC}_0$ *of* $\mathcal{C}_0$ *by* $\mathcal{AC}_0$*. In particular,* $\mathcal{C}_\infty$ *is triangulated.*

Remark 2.40. (a) The collection S of all r-isomorphisms (for every $r \geq 0$) forms a multiplicative system in $\mathcal{C}_0$. The Verdier quotient above is the same as the localization $S^{-1}\mathcal{C}_0$ of $\mathcal{C}_0$ by S. For a definition of the localization, see [40, Subsection 1.6].

(b) A result somewhat similar to Proposition 2.39, established for Tamarkin categories, appears in [37, Proposition 6.7].

Proof. For the first point of the proposition we first notice that the subcategory of acyclics, $\mathcal{AC}$, is a persistence category. It is obviously endowed with a shift functor by restricting the shift functor of $\mathcal{C}$. Moreover, its 0-level is clearly a full subcategory of $\mathcal{C}_0$. Finally, Lemma 2.29 implies that $\mathcal{AC}_0$ is a triangulated subcategory of $\mathcal{C}_0$, which implies that it is a TPC.

We now turn to the second point of the proposition. By inspecting the definition of Verdier localization (for instance in [43, Chapter 2]), we see that the localization of $\mathcal{C}_0$ at $\mathcal{AC}_0$, denoted by $\mathcal{C}_0/\mathcal{AC}_0$, is a category with the same objects as $\mathcal{C}_0$ and having as morphisms $A \to B$ equivalence classes of roof diagrams:

$$A \xleftarrow{u} A' \xrightarrow{f'} B$$

with u an r-isomorphism for some $r \geq 0$, and $f' \in \hom_{\mathcal{C}_0}(A', B)$. Two roof diagrams $A \xleftarrow{u} A' \xrightarrow{f'} B$ and $A \xleftarrow{u_1} A'_1 \xrightarrow{f'_1} B$ are equivalent if they are related by a third roof

diagram $A \xleftarrow{u_2} A'_2 \xrightarrow{f'_2} B$ in the sense that there are maps a', a'_1 in $\mathcal{C}_0$ making the following diagram commutative:

$$\begin{array}{ccccc}
A & \xleftarrow{u} & A' & \xrightarrow{f'} & B \\
\uparrow{\scriptstyle \mathbb{1}_A} & & \uparrow{\scriptstyle a'} & & \uparrow{\scriptstyle \mathbb{1}_B} \\
A & \xleftarrow{u_2} & A'_2 & \xrightarrow{f'_2} & B \\
\downarrow{\scriptstyle \mathbb{1}_A} & & \downarrow{\scriptstyle a'_1} & & \downarrow{\scriptstyle \mathbb{1}_B} \\
A & \xleftarrow{u_1} & A'_1 & \xrightarrow{f'_1} & B
\end{array}$$

The category $\mathcal{C}_\infty$ appears in Definition 2.9. Its objects are the same as those of $\mathcal{C}_0$ and its morphisms are

$$\hom_{\mathcal{C}_\infty}(A, B) = \varinjlim_{r\to\infty} \hom^r_{\mathcal{C}}(A, B) = \varinjlim_{r\to\infty} \hom_{\mathcal{C}_0}(\Sigma^r A, B),$$

where the second equality comes from (2.9). Given a morphism $f \in \hom_{\mathcal{C}_\infty}(A, B)$, this means that f is represented by $\bar{f} : \Sigma^r A \to B$ for some $r \geq 0$ as well as by all compositions $\Sigma^s A \xrightarrow{\eta_{s-r}} \Sigma^r A \xrightarrow{\bar{f}} B$. We now define a functor

$$\Phi : \mathcal{C}_\infty \to \mathcal{C}_0/\mathcal{A}\mathcal{C}_0.$$

It is the identity on objects and for a morphism $f \in \hom_{\mathcal{C}_\infty}(A, B)$ we let $\Phi(f)$ be the equivalence class of a roof diagram $A \xleftarrow{\eta_r} \Sigma^r A \xrightarrow{\bar{f}} B$, where $\bar{f} : \Sigma^r A \to B$ represents f. Any two roof diagrams that are associated to two representatives of f are immediately seen to be equivalent, and as a result Φ is well defined.

It remains to show that Φ is an isomorphism. Surjectivity is immediate. Fix H a roof diagram $A \xleftarrow{a} A' \xrightarrow{b} B$. As a is an r-isomorphism (for some r), we deduce from Proposition 2.28 the existence of a right r-inverse $\phi : \Sigma^r A \to A'$ of a such that $a \circ \phi = \eta^A_r$. Therefore, we may define a new roof diagram H', $A \xleftarrow{\eta_r} \Sigma^r A \xrightarrow{b\circ\phi} B$. The roof diagrams H' and H are clearly equivalent, and thus their equivalence class belongs to the image of Φ.

We now show that Φ is injective. For this we consider the commutative diagram

$$\begin{array}{ccccc}
A & \xleftarrow{\eta^A_{r'}} & \Sigma^{r'} A & \xrightarrow{f''} & B \\
\uparrow{\scriptstyle \mathbb{1}_A} & & \uparrow{\scriptstyle a''} & & \uparrow{\scriptstyle \mathbb{1}_B} \\
A & \xleftarrow{u} & A' & \xrightarrow{f} & B \\
\downarrow{\scriptstyle \mathbb{1}_A} & & \downarrow{\scriptstyle a'} & & \downarrow{\scriptstyle \mathbb{1}_B} \\
A & \xleftarrow{\eta^A_r} & \Sigma^r A & \xrightarrow{f'} & B
\end{array}$$

with each row a roof diagram. We need to show that f'' and f' represent the same element in $\mathcal{C}_\infty$. In fact, u is an s-isomorphism for some $s \geq 0$. Let $v : \Sigma^s A \to A'$ be a right inverse of u. Now consider the commutative diagram

$$\begin{array}{ccccc}
A & \xleftarrow{\eta^A_{r'}} & \Sigma^{r'} A & \xrightarrow{f''} & B \\
\uparrow{\scriptstyle \mathbb{1}_A} & & \uparrow{\scriptstyle \bar{a}''} & & \uparrow{\scriptstyle \mathbb{1}_B} \\
A & \xleftarrow{\eta_s} & \Sigma^s A & \xrightarrow{\bar{f}} & B \\
\downarrow{\scriptstyle \mathbb{1}_A} & & \downarrow{\scriptstyle \bar{a}'} & & \downarrow{\scriptstyle \mathbb{1}_B} \\
A & \xleftarrow{\eta^A_r} & \Sigma^r A & \xrightarrow{f'} & B
\end{array}$$

where $\bar{a}' = a' \circ v$, $\bar{a}'' = a'' \circ v$, $\bar{f} = f \circ v$. Notice that we have $\eta^A_r \circ \bar{a}' = \eta^A_s = \eta_r \circ \eta^A_{s-r}$. This means, by Corollary 2.35, that $\bar{a}' \simeq_r \eta^A_{s-r}$. Similarly, we have $\bar{a}'' \simeq_{r'} \eta^A_{s-r'}$. For $r'' \geq \max\{r, r'\}$ we deduce $\bar{a}' \circ \eta^A_{r''} = \eta^A_{s-r+r''}$ and $\bar{a}'' \circ \eta^A_{r''} = \eta^A_{s-r'+r''}$. Thus, by composing on the middle node with $\eta^A_{r''} : \Sigma^{s+r''} A \to \Sigma^s A$, we get a new commutative diagram similar to the one above but with $\bar{a}''$ and $\bar{a}'$ being replaced with $\eta^A_{s-r+r''}$ and $\eta^A_{s-r+r'}$, respectively. We deduce that f'' and f' give in $\mathcal{C}_\infty$ the same element as $\bar{f} \circ \eta^A_{r''}$, which concludes the proof. ∎

Remark 2.41. By the properties of Verdier localization, the category $\mathcal{C}_\infty$ is triangulated in such a way that, by definition, a triangle in $\mathcal{C}_\infty$ is exact if it is isomorphic with the image (in $\mathcal{C}_\infty$) of an exact triangle from $\mathcal{C}_0$.

2.3.3 Weighted exact triangles

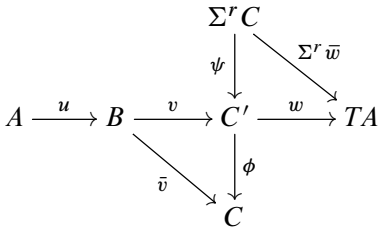
The key feature of a triangulated persistence category $\mathcal{C}$ is that there is a natural way to associate weights to a class of triangles larger than the exact triangles in $\mathcal{C}_0$.

Definition 2.42. A *strict exact triangle in* $\mathcal{C}$ is a pair $\widetilde{\Delta} = (\Delta, r)$ where $r \in [0, +\infty)$ and Δ is a diagram

$$\Delta : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r} TA$$

in $\mathcal{C}_0$ with $\bar{u} \in \hom^0_{\mathcal{C}}(A, B)$, $\bar{v} \in \hom^0_{\mathcal{C}}(B, C)$, and $\bar{w} \in \hom^0_{\mathcal{C}}(C, \Sigma^{-r} TA)$, such that the following holds. There exists an exact triangle $A \xrightarrow{u} B \xrightarrow{v} C' \xrightarrow{w} TA$ in $\mathcal{C}_0$, with $u = \bar{u}$, an r-isomorphism $\phi : C' \to C$, and a right r-inverse of ϕ, denoted by

$\psi : \Sigma^r C \to C'$, such that the following diagram commutes:



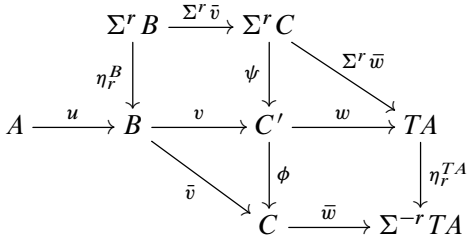
(2.20)

The weight of the strict exact triangle $\widetilde{\Delta}$ is the number r, and we denote it by $w(\widetilde{\Delta})$.

Remark 2.43. (a) To simplify terminology, we will often denote strict exact triangles by the diagram Δ, with the weight identified implicitly by the amount of down "shift" of the last term. Notice that if $\Sigma^s A \neq A$ for all s, then the diagram Δ determines the weight of the triangle. However, when this is not the case, it is necessary to indicate the weight explicitly. For example, for any $r \geq 0$, the pair $(0 \to X \xrightarrow{\mathbb{1}_X} X \to 0, r)$ is a strict exact triangle of weight r because $\Sigma^s 0 = 0$ in $\mathcal{C}_0$ and two such triangles are different as soon as the corresponding weights are different. In what follows, we will not always write strict exact triangles as pairs. We will often simply write that a diagram Δ as above is strict exact of weight $w(\Delta) = r$. Although there is a slight imprecision in writing $w(\Delta) = r$ (since Δ does not determine r), the meaning of this should be clear: (Δ, r) is a strict triangle of weight r.

(b) Any exact triangle in $\mathcal{C}_0$ is a strict exact triangle of weight 0. Conversely, it is a simple exercise to see that a strict exact triangle of weight 0 is exact as a triangle in $\mathcal{C}_0$.

(c) Consider the following diagram:



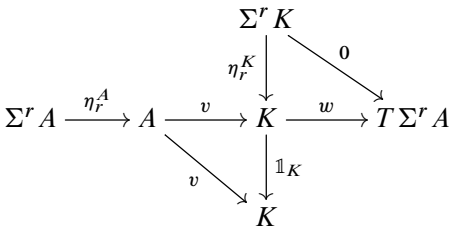
which is derived from the commutative diagram (2.20). The two squares in the diagram are not commutative in general, but they are r-commutative. Indeed, since ϕ is an r-isomorphism, let $\widetilde{\psi} : C \to \Sigma^{-r} C'$ be a left r-inverse of ϕ. As ψ is a right r-inverse of ϕ, we deduce from Proposition 2.28 (i) that $\Sigma^{-r}\psi \simeq_r \widetilde{\psi}$. Therefore, $\bar{w} \circ \phi = \Sigma^{-r} w \circ \Sigma^{-r}\psi \circ \phi \simeq_r \Sigma^{-r} w \circ \widetilde{\psi} \circ \phi = \Sigma^{-r} w \circ \eta_r^{C'} = \eta_r^{TA} \circ w$. Using Corollary 2.35, we also see that $\psi \circ \Sigma^r \bar{v} \simeq_r v \circ \eta_r^B$, because $\phi \circ \psi \circ \Sigma^r \bar{v} = \phi \circ v \circ \eta_r^B$.

(d) Because T commutes with Σ and with the natural transformations η, it immediately follows that this functor preserves strict exact triangles as well as their weight.

Example 2.44. Recall that the map $\Sigma^r A \xrightarrow{\eta_r^A} A$ embeds into an exact triangle in $\mathcal{C}_0$, $\Sigma^r A \xrightarrow{\eta_r^A} A \xrightarrow{v} K \xrightarrow{w} T\Sigma^r A$, where K is r-acyclic. We claim that the diagram

$$\Sigma^r A \xrightarrow{\eta_r^A} A \xrightarrow{v} K \xrightarrow{0} TA$$

is a strict exact triangle of weight r. Indeed, we have the commutative diagram



where the upper-right triangle is commutative since K is r-acyclic (so $\eta_r^K = 0$ by Lemma 2.18 (i)). Moreover, $\mathbb{1}_K$ is an r-isomorphism (recall that $T\Sigma^r A = \Sigma^r TA$).

Note that the diagram

$$\Sigma^r A \xrightarrow{\eta_r^A} A \to 0 \to TA$$

is also a strict exact triangle of weight r.

Proposition 2.45 (Weight invariance). *Strict exact triangles satisfy the following two properties:*

(i) *Suppose the two diagrams $A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C$ and $A' \xrightarrow{\bar{u}'} B' \xrightarrow{\bar{v}'} C'$ are isomorphic in $\mathcal{C}_0$, i.e., we have the following commutative diagram in $\mathcal{C}_0$:*

$$\begin{array}{ccccc} A & \xrightarrow{\bar{u}} & B & \xrightarrow{\bar{v}} & C \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle h} \\ A' & \xrightarrow{\bar{u}'} & B' & \xrightarrow{\bar{v}'} & C' \end{array} \tag{2.21}$$

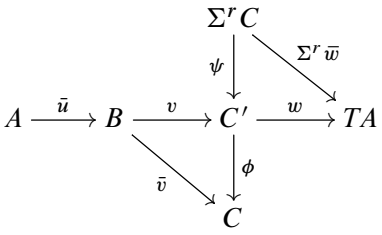
Then $A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C$ completes to a strict exact triangle of weight r, denoted by $\Delta : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r} TA$ if and only if $A' \xrightarrow{\bar{u}'} B' \xrightarrow{\bar{v}'} C'$ completes to a strict exact triangle of weight r, denoted by $\Delta' : A' \xrightarrow{\bar{u}'} B' \xrightarrow{\bar{v}'} C' \xrightarrow{\bar{w}'} \Sigma^{-r} TA'$. Moreover, Δ and Δ' are isomorphic in $\mathcal{C}_0$.

(ii) *If $\Delta : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r} TA$ satisfies $w(\Delta) = r$, then $\Delta' : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}'} \Sigma^{-r-s} TA$ satisfies $w(\Delta') = r + s$ for $s \geq 0$, where $\bar{w}'$ is the composition*

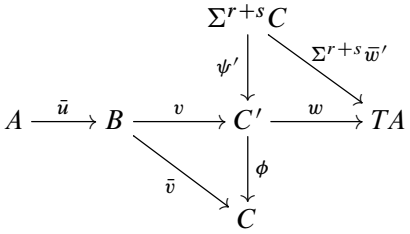
$$C \xrightarrow{\bar{w}} \Sigma^{-r} TA \xrightarrow{\eta_s^{\Sigma^{-r} TA}} \Sigma^{-r-s} TA.$$

Proof. (i) The property claimed here immediately follows from the fact that, in $\mathcal{C}_0$, all 0-isomorphisms admit inverses. Therefore, if $A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C$ completes to a strict exact triangle Δ, then the desired map $\bar{w}'$ for Δ' can be chosen as $\bar{w}' = T(f) \circ \bar{w} \circ h^{-1}$, where h^{-1} is the inverse of map h in (2.21). The weight can be easily deduced from Definition 2.42.

(ii) By definition, there exists a commutative diagram



Set $\psi' = \psi \circ \eta_s^{\Sigma^r C} \in \hom^r(\Sigma^{r+s} C, C')$. Then ψ' is an $(r+s)$-right inverse of ϕ. Consider the diagram



where $\bar{w}' : C \to \Sigma^{-r-s} TA$ is defined by $\bar{w}' = \eta_s^{\Sigma^{-r} TA} \circ \bar{w}$, and notice that the upper-right triangle is commutative. ∎

Proposition 2.46 (Weighted rotation property). *Given a strict exact triangle*

$$\Delta : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r} TA$$

satisfying $w(\Delta) = r$, there exists a triangle

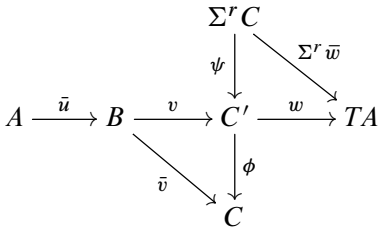
$$R(\Delta) : B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}'} \Sigma^{-r} TA \xrightarrow{-\bar{u}'} \Sigma^{-2r} TB$$

satisfying $w(R(\Delta)) = 2r$, where $\bar{w}' \simeq_r \bar{w}$ and $\bar{u}'$ is the composition

$$\bar{u}' : \Sigma^{-r} TA \xrightarrow{\Sigma^{-r} T\bar{u}} \Sigma^{-r} TB \xrightarrow{\eta_r^{\Sigma^{-r} TB}} \Sigma^{-2r} TB.$$

We call $R(\Delta)$ the (first) positive rotation of Δ.

Proof. By definition, there exists a commutative diagram



where $A \xrightarrow{\bar{u}} B \xrightarrow{v} C' \xrightarrow{w} TA$ is an exact triangle in $\mathcal{C}_0$ and ψ is a right r-inverse of ϕ. By the rotation property of $\mathcal{C}_0$, $B \xrightarrow{v} C' \xrightarrow{w} TA \xrightarrow{-T\bar{u}} TB$ is an exact triangle in $\mathcal{C}_0$. We now construct the following diagram in $\mathcal{C}_0$ in which the upper squares will be commutative and the lower square r-commutative:

$$\begin{array}{ccccccc}
B & \xrightarrow{v} & C' & \xrightarrow{w} & TA & \xrightarrow{-T\bar{u}} & TB \\
{\scriptstyle \mathbb{1}_B}\downarrow & & \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \bar{\phi}} & & \downarrow{\scriptstyle \mathbb{1}_{TB}} \\
B & \xrightarrow{\bar{v}} & C & \xrightarrow{w'} & TA' & \xrightarrow{u'} & TB \\
 & & {\scriptstyle \Sigma^{-r}\psi}\downarrow & & \downarrow{\scriptstyle \bar{\psi}} & & \\
 & & \Sigma^{-r}C' & \xrightarrow{\Sigma^{-r}w} & \Sigma^{-r}TA & &
\end{array} \tag{2.22}$$

Here the second row of maps comes from embedding $B \xrightarrow{\bar{v}} C$ into an exact triangle $B \xrightarrow{\bar{v}} C \xrightarrow{w'} A'' \to TB$ for some A'' in $\mathcal{C}_0$ and $A' = T^{-1}A''$. The map $\bar{\phi}$ is then induced by the functoriality of triangles in $\mathcal{C}_0$ and is an r-isomorphism by Proposition 2.31 (i). So far this gives the three upper squares of the diagram and their commutativity. To construct the lower square, let $\bar{\psi}$ be a left r-inverse of $\bar{\phi}$ (i.e., $\bar{\psi} \circ \bar{\phi} = \eta_r^{TA}$). By Corollary 2.33, $\bar{\psi}$ is a $2r$-isomorphism.

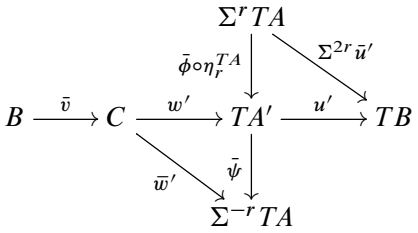
We claim that the lower square in the diagram (2.22) is r-commutative, and therefore we have $\bar{\psi} \circ w' \simeq_r \Sigma^{-r} w \circ \Sigma^{-r}\psi = \overline{w}$.

Indeed, let ψ' be a left r-inverse of ϕ. By using the commutativity of the upper-middle square in the diagram (2.22), we deduce $\bar{\psi} \circ w' \circ \phi = \Sigma^{-r} w \circ \psi' \circ \phi$. As ϕ is an r-isomorphism, we obtain

$$\bar{\psi} \circ w' \simeq_r \Sigma^{-r} w \circ \psi' \simeq_r \Sigma^{-r} w \circ \Sigma^{-r}\psi = \overline{w},$$

because, by Proposition 2.28, we have $\Sigma^{-r}\psi \simeq_r \psi'$. This shows that the lower square is r-commutative and that the related r-identity holds.

We next consider the diagram



where $\bar{u}' = \Sigma^{-2r}(u' \circ \bar{\phi} \circ \eta_r^{TA})$ and $\bar{w}' = \bar{\psi} \circ w'$. Given that $\bar{\psi}$ is a $2r$-isomorphism, this means that we have a strict exact triangle of weight $2r$ of the form

$$B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}'} TA' \xrightarrow{\bar{u}'} \Sigma^{-2r} TB.$$

We already know that $\bar{w}' = \bar{\psi} \circ w' \simeq_r \bar{w}$. On the other hand,

$$\Sigma^{2r}\bar{u}' = u' \circ (\bar{\phi} \circ \eta_r^{TA}) = -T\bar{u} \circ \eta_r^{TA} = -\eta_r^{\Sigma^r TB} \circ \Sigma^r T\bar{u},$$

which concludes the proof. ∎

Remark 2.47. An entirely similar argument also shows that there exists a strict exact triangle of weight $2r$ and of the form

$$R^{-1}(\Delta) : T^{-1}\Sigma^r C \to A \to B \to \Sigma^{-r} C,$$

which is the (first) negative rotation of Δ. Note that $R^{-1}(R(\Delta)) \neq R(R^{-1}(\Delta)) \neq \Delta$.

Remark 2.48. Proposition 2.46 describes a rotation of weighted exact triangles that does *not* preserve weights. Indeed, the rotation of the weight-r triangle Δ considered there has weight $2r$. It is not clear to what extent one can improve this. Ideally, one would like to be able to rotate Δ into a weighted exact triangle $B \to C \to \Sigma^{-r} TA \to \Sigma^{-r} TB$ of the *same* weight r. There is some evidence, coming from symplectic topology, indicating that in certain circumstances this might be possible (see Section 3.5.1.9). However, the algebraic setting in this memoir, in particular the definition of weighted exact triangles, might be too general to render this feasible, at least without additional assumptions on Δ.

Proposition 2.49 (Weighted octahedral formula). *Given two strict exact triangles*

$$\Delta_1 : E \xrightarrow{\beta} F \xrightarrow{\alpha} X \xrightarrow{k} \Sigma^{-r} TE$$

and

$$\Delta_2 : X \xrightarrow{u} A \xrightarrow{\gamma} B \xrightarrow{b} \Sigma^{-s} TX$$

with $w(\Delta_1) = r$ and $w(\Delta_2) = s$, there exists a diagram

$$\begin{array}{ccccccc}
E & \longrightarrow & 0 & \longrightarrow & TE & \longrightarrow & TE \\
{\scriptstyle\beta}\downarrow & & \downarrow & & \downarrow & & \downarrow{\scriptstyle T\beta} \\
F & \xrightarrow{u\circ\alpha} & A & \longrightarrow & C & \longrightarrow & TF \\
{\scriptstyle\alpha}\downarrow & & {\scriptstyle\mathbb{1}}\downarrow & & \downarrow & & \downarrow{\scriptstyle(\Sigma^{-s}T\alpha)\circ\eta_s^{TF}} \\
X & \xrightarrow{u} & A & \xrightarrow{\gamma} & B & \xrightarrow{b} & \Sigma^{-s}TX \\
{\scriptstyle k}\downarrow & & \downarrow & & \downarrow & & \downarrow{\scriptstyle\Sigma^{-s}(Tk)} \\
\Sigma^{-r}TE & \longrightarrow & 0 & \longrightarrow & \Sigma^{-r-s}T^2E & \longrightarrow & \Sigma^{-r-s}T^2E
\end{array} \tag{2.23}$$

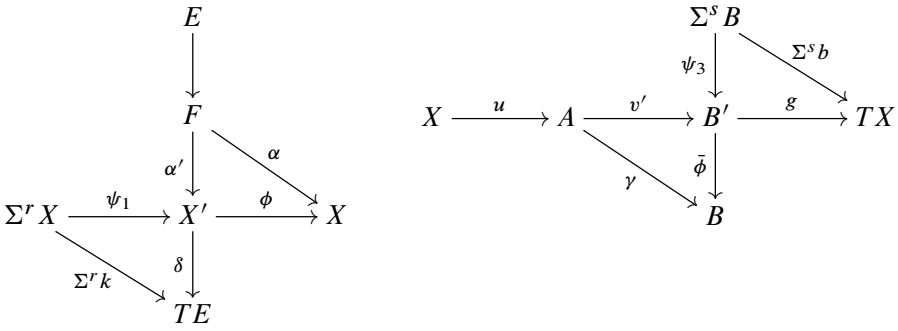
with all squares commutative except for the bottom-right one, which is r-anti-commutative, such that the triangles

$$\Delta_3 : F \to A \to C \to TF \quad \textit{and} \quad \Delta_4 : TE \to C \to B \to \Sigma^{-r-s}T^2E$$

are strict exact with $w(\Delta_3) = 0$ and $w(\Delta_4) = r + s$.

By forgetting the Σ's (or assuming that $r = s = 0$), this is equivalent to the usual octahedral axiom in a triangulated category (namely $\mathcal{C}_0$) and the bottom-right square is commutative up to sign (or anti-commutative).

Proof of Proposition 2.49. By definition, there are two commutative diagrams



(2.24)

with ϕ an r-isomorphism and $\bar{\phi}$ an s-isomorphism, and ψ_1 and ψ_3 are their right r- and s-inverses, respectively. By the octahedral axiom in $\mathcal{C}_0$, we construct the following diagram, with all squares commutative except for the bottom-right one, which

is anti-commutative:

$$\begin{array}{ccccccc}
E & \longrightarrow & 0 & \longrightarrow & TE & \longrightarrow & TE \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
F & \xrightarrow{\alpha''} & A & \longrightarrow & C & \longrightarrow & TF \\
{\scriptstyle\alpha'}\downarrow & & {\scriptstyle\mathbb{1}_A}\downarrow & & {\scriptstyle w}\downarrow & & {\scriptstyle T\alpha'}\downarrow \\
X' & \xrightarrow{u\circ\phi} & A & \xrightarrow{v} & B'' & \xrightarrow{\theta} & TX' \\
{\scriptstyle\delta}\downarrow & & \downarrow & & {\scriptstyle t}\downarrow & & {\scriptstyle T\delta}\downarrow \\
TE & \longrightarrow & 0 & \longrightarrow & T^2E & \xrightarrow{\mathbb{1}_{T^2E}} & T^2E
\end{array}$$

Thus $\alpha'' = u \circ \phi \circ \alpha'$, $t = -T\delta \circ \theta$. We denote by

$$\Delta_3 : F \xrightarrow{\alpha''} A \to C \to TF$$

the respective exact triangle in $\mathcal{C}_0$ so that, as in Remark 2.43, $w(\Delta_3) = 0$. The map $w \in \hom^0(C, B'')$ is induced from the commutativity of the middle-left triangle. We now consider the following diagram:

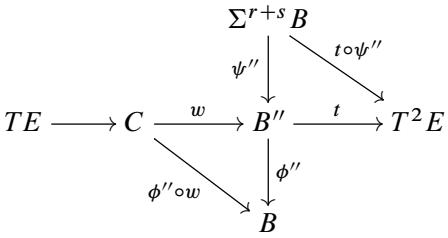
$$\begin{array}{ccccccc}
F & \longrightarrow & A & \longrightarrow & C & \longrightarrow & TF \\
{\scriptstyle\alpha'}\downarrow & & {\scriptstyle\mathbb{1}_A}\downarrow & & {\scriptstyle w}\downarrow & & {\scriptstyle T\alpha'}\downarrow \\
X' & \xrightarrow{u\circ\phi} & A & \xrightarrow{v} & B'' & \xrightarrow{\theta} & TX' \\
{\scriptstyle\phi}\downarrow & & {\scriptstyle\mathbb{1}_A}\downarrow & & {\scriptstyle\phi'}\downarrow & & {\scriptstyle T\phi}\downarrow \\
X & \xrightarrow{u} & A & \xrightarrow{v'} & B' & \xrightarrow{g} & TX \\
 & & {\scriptstyle\mathbb{1}_A}\downarrow & & {\scriptstyle\bar{\phi}}\downarrow & & \\
 & & A & \xrightarrow{\gamma} & B & \xrightarrow{b} & \Sigma^{-s}TX
\end{array} \tag{2.25}$$

The three long rows are exact triangles in $\mathcal{C}_0$, and we deduce the existence of $\phi' \in \hom^0(B'', B')$ making the adjacent squares commutative. This is an r-isomorphism by Proposition 2.31 (i). We fix a right r-inverse $\psi_2 \in \hom^0(\Sigma^r B', B'')$ of ϕ'. The composition $\phi'' = \bar{\phi} \circ \phi' \in \hom^0(B'', B)$ is an $(r+s)$-isomorphism by Proposition 2.28 (iii). Let $\psi'' = \psi_2 \circ \Sigma^r \psi_3$ (recall ψ_3 from (2.24)) and notice that ψ'' is a right $(r+s)$-inverse of ϕ''.

We are now able to define the triangle Δ_4:

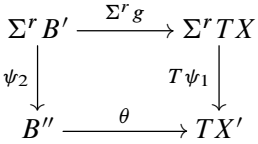
$$\Delta_4 : TE \to C \xrightarrow{\phi''\circ w} B \xrightarrow{\Sigma^{-r-s}(t\circ\psi'')} \Sigma^{-r-s}T^2E.$$

The following commutative diagram shows that Δ_4 is strict exact and $w(\Delta_4) = r + s$:

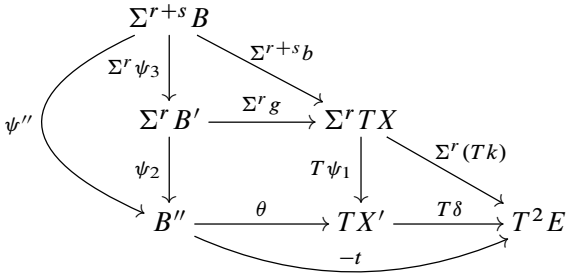


It is easy to check that all the squares in (2.23), except the bottom-right one, are commutative.

We now check the r-anti-commutativity of the bottom-right square. We need to show that $\Sigma^{-s}(Tk) \circ b \simeq_r -\Sigma^{-r-s}(t \circ \psi'')$, which is equivalent to $\Sigma^r(Tk) \circ \Sigma^{r+s}b \simeq_r -t \circ \psi''$. Given that the $B''(TX')(TX)B'$ square in (2.25) commutes, Corollary 2.38 yields the following r-commutative diagram:



Now consider the following diagram, commutative except for the middle square, which is r-commutative:



Write

$$\begin{aligned} -t \circ \psi'' &= (T\delta) \circ (\theta \circ \psi_2) \circ \Sigma^r \psi_3 \\ &\simeq_r (T\delta) \circ ((T\psi_1) \circ \Sigma^r g) \circ \Sigma^r \psi_3 = \Sigma^r(Tk) \circ \Sigma^{r+s} b, \end{aligned}$$

which completes the proof. ∎

Given a triple of maps $\Delta : A \xrightarrow{u} B \xrightarrow{v} C \xrightarrow{w} D$ with shifts $\lceil u \rceil, \lceil v \rceil, \lceil w \rceil \in \mathbb{R}$, it is useful to introduce a special notation for an associated triple in $\mathcal{C}_0$, denoted by

$\Sigma^{s_1,s_2,s_3,s_4}\Delta$, for $s_1, s_2, s_3, s_4 \in \mathbb{R}$ satisfying the following relations:

$$\begin{aligned} -s_1 + s_2 + \lceil u \rceil &\leq 0, \\ -s_2 + s_3 + \lceil v \rceil &\leq 0, \\ -s_3 + s_4 + \lceil w \rceil &\leq 0. \end{aligned}$$

The triple $\Sigma^{s_1,s_2,s_3,s_4}\Delta$ has the form

$$\Sigma^{s_1} A \xrightarrow{\bar{u}} \Sigma^{s_2} B \xrightarrow{\bar{v}} \Sigma^{s_3} \xrightarrow{\bar{w}} \Sigma^{s_4} D, \tag{2.26}$$

where $\bar{u}$ is the composition of the composition $\Sigma^{s_1} A \xrightarrow{(\eta_{s_1,0})_A} A \xrightarrow{u} B \xrightarrow{(\eta_{0,s_2})_B} \Sigma^{s_2} B$ and the persistence structure map $i_{-s_1+s_2+\lceil u \rceil,0}$, i.e.,

$$\bar{u} = i_{-s_1+s_2+\lceil u \rceil,0}\big((\eta_{0,s_2})_B \circ u \circ (\eta_{s_1,0})_A\big).$$

The definitions of $\bar{v}$ and $\bar{w}$ are similar and, in particular, $\lceil \bar{u} \rceil = \lceil \bar{v} \rceil = \lceil \bar{w} \rceil = 0$. The inequalities above ensure that the resulting triangle (2.26) has all morphisms in $\mathcal{C}_0$.

For $s_1 = s_2 = s_3 = s_4 = k$ (which implies that $\lceil u \rceil, \lceil v \rceil, \lceil w \rceil \leq 0$) we write, for brevity,

$$\Sigma^k \Delta := \Sigma^{s_1,s_2,s_3,s_4}\Delta.$$

Remark 2.50. Assume that $\Delta : A \xrightarrow{u} B \xrightarrow{v} C \xrightarrow{w} \Sigma^{-r} TA$ is strict exact of weight $w(\Delta) = r$.

(a) It is a simple exercise to show that the triangle $\Sigma^k \Delta : \Sigma^k A \to \Sigma^k B \to \Sigma^k C \to \Sigma^{-r+k} TA$ is strict exact and $w\big(\Sigma^k \Delta\big) = r$.

(b) For $s \geq 0$, Proposition 2.45 (ii) claims that $\Sigma^{0,0,0,-s}\Delta$ is strict exact of weight $r + s$. It is again an easy exercise to see that $\Sigma^{0,0,-s,-s}\Delta$ is strict exact of weight $r + s$.

Proposition 2.51 (Functoriality of triangles). *Consider two strict exact triangles as below, with $f \in \hom^0(A_1, A_2)$ and $g \in \hom^0(B_1, B_2)$:*

$$\begin{array}{lccccccc} \Delta_1 : & A_1 & \xrightarrow{\bar{u}_1} & B_1 & \xrightarrow{\bar{v}_1} & C_1 & \xrightarrow{\bar{w}_1} & \Sigma^{-r} TA_1 \\ & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & & & \\ \Delta_2 : & A_2 & \xrightarrow{\bar{u}_2} & B_2 & \xrightarrow{\bar{v}_2} & C_2 & \xrightarrow{\bar{w}_2} & \Sigma^{-s} TA_2 \end{array}$$

and assume that $w(\Delta_1) = r$ and $w(\Delta_2) = s$. Then there exists a morphism $h \in \hom^0(C_1, \Sigma^{-r} C_2)$ inducing maps relating the triangles $\Delta_1 \to \Sigma^{0,0,-r,-r}\Delta_2$, as in

the following diagram:

$$\begin{array}{ccccccc} A_1 & \longrightarrow & B_1 & \longrightarrow & C_1 & \longrightarrow & \Sigma^{-r}TA_1 \\ {\scriptstyle f}\downarrow & & {\scriptstyle g}\downarrow & & {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle \eta_s^{TA_2}\circ\Sigma^{-r}Tf} \\ A_2 & \longrightarrow & B_2 & \longrightarrow & \Sigma^{-r}C_2 & \longrightarrow & \Sigma^{-r-s}TA_2 \end{array}$$

where the middle square is r-commutative and the right square is s-commutative.

The proof is left as an exercise.

Proposition 2.52. *Let $\Delta : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r}TA$ be a strict exact triangle of weight r in $\mathcal{C}$ and let $\Delta' : A \xrightarrow{\bar{u}} B \xrightarrow{v} C' \xrightarrow{w} TA$ be the exact triangle in $\mathcal{C}_0$ associated to Δ as in Definition 2.42. There are morphisms of triangles $h : \Sigma^{2r}\Delta \to \Delta'$ and $h' : \Delta' \to \Sigma^{-r}\Delta$ such that the compositions $h' \circ h$ and $h \circ \Sigma^{3r}h'$ have as vertical maps the shifted natural transformations $\eta_{3r}^{(-)}$ defined in (2.14).*

Proof. We use the notation in Definition 2.42 and consider the diagram below:

$$\begin{array}{ccccccc} \Sigma^r A & \xrightarrow{\Sigma^r\bar{u}} & \Sigma^r B & \xrightarrow{\Sigma^r\bar{v}} & \Sigma^r C & \xrightarrow{\Sigma^r\bar{w}} & TA \\ {\scriptstyle \eta_r^A}\downarrow & & {\scriptstyle \eta_r^B}\downarrow & & {\scriptstyle \psi}\downarrow & & \downarrow{\scriptstyle \mathbb{1}_{TA}} \\ A & \xrightarrow{\bar{u}} & B & \xrightarrow{v} & C' & \xrightarrow{w} & TA \\ {\scriptstyle \mathbb{1}_A}\downarrow & & {\scriptstyle \mathbb{1}_B}\downarrow & & {\scriptstyle \phi}\downarrow & & \downarrow{\scriptstyle \eta_r^{TA}} \\ A & \xrightarrow{\bar{u}} & B & \xrightarrow{\bar{v}} & C & \xrightarrow{\bar{w}} & \Sigma^{-r}TA \end{array}$$

Set $h_1 = (\eta_r^A, \eta_r^B, \psi, \mathbb{1}_{TA})$ and $h_1' = (\mathbb{1}_A, \mathbb{1}_B, \phi, \eta_r^{TA})$. Notice that h_1 as well as h_1' are not morphisms of triangles because the bottom-right square is only r-commutative, and the same is true for the top-middle square, as discussed in Remark 2.43 (c). Let $h = h_1 \circ \eta_r$ and $h' = \eta_r \circ h_1'$ (where we view η_r as a quadruple of morphisms of the form η_r^-). It follows that both h and h' are morphisms of triangles. Moreover, given that $\phi \circ \psi = \eta_r^C$, it is clear that $h' \circ h = \eta_{3r}$. The other composition, $h \circ \Sigma^{3r}h'$, has a term of the form $\psi \circ \eta_r \circ \eta_r \circ \Sigma^{3r}$; thus, by Proposition 2.28, this coincides with $\eta_{3r}^{C'}$, as claimed. ∎

Remark 2.53. Proposition 2.52 shows that a strict exact triangle of weight r is approximately isomorphic, in a sense similar to interleaving, to an exact triangle in $\mathcal{C}_0$.

2.3.4 Fragmentation pseudometrics on Obj($\mathcal{C}$)

In a triangulated persistence category there is a natural notion of iterated cone decomposition, similar to the corresponding notion in the triangulated setting from Section 2.1.

Definition 2.54. Let $\mathcal{C}$ be a triangulated persistence category, and let $X \in \mathrm{Obj}(\mathcal{C})$. An *iterated cone decomposition* D of X with *linearization* $(X_1, X_2, \ldots, X_n)$, where $X_i \in \mathrm{Obj}(\mathcal{C})$, consists of a family of strict exact triangles in $\mathcal{C}$,

$$\begin{cases} \Delta_1 : X_1 \to 0 \to Y_1 \to \Sigma^{-r_1} T X_1 \\ \Delta_2 : X_2 \to Y_1 \to Y_2 \to \Sigma^{-r_2} T X_2 \\ \Delta_3 : X_3 \to Y_2 \to Y_3 \to \Sigma^{-r_3} T X_3 \\ \qquad\qquad \vdots \\ \Delta_n : X_n \to Y_{n-1} \to X \to \Sigma^{-r_n} T X_n. \end{cases}$$

The weight of such a cone decomposition is defined by

$$w(D) = \sum_{i=1}^{n} w(\Delta_i).$$

The linearization of D is denoted by $\ell(D) = (X_1, \ldots, X_n)$.

Proposition 2.55. *Assume that X admits an iterated cone decomposition D with linearization $(X_1, \ldots, X_n)$ and, for some $i \in \{1, \ldots, n\}$, X_i admits an iterated cone decomposition D' with linearization $(A_1, \ldots, A_k)$. Then X admits an iterated cone decomposition D'' with linearization*

$$(X_1, \ldots, X_{i-1}, TA_1, \ldots, TA_k, X_{i+1}, \ldots, X_n).$$

Moreover, the weights of these cone decompositions satisfy $w(D'') = w(D) + w(D')$.

A cone decomposition D'' as in the statement of Proposition 2.55 is called a *refinement* of the cone decomposition D with respect to D'.

Example 2.56. A single strict exact triangle $A \to B \to X \to \Sigma^{-r} TA$ can be regarded as a cone decomposition D of X with linearization $(T^{-1}B, A)$ such that $w(D) = r$. Assume that A fits into a second strict exact triangle $E \to F \to A \to \Sigma^{-s} TE$ of weight s. Thus we have a cone decomposition D' of A, with linearization $(T^{-1}F, E)$ and $w(D') = s$. Diagram (2.23) from Proposition 2.49 yields the commutative diagram

$$\begin{array}{ccccc} E & \longrightarrow & 0 & \longrightarrow & TE \\ \downarrow & & \downarrow & & \downarrow \\ F & \longrightarrow & B & \longrightarrow & Y \\ \downarrow & & \downarrow & & \downarrow \\ A & \longrightarrow & B & \longrightarrow & X \end{array}$$

for some object $Y \in \mathrm{Obj}(\mathcal{C})$. In particular, we obtain a strict exact triangle $TE \to Y \to X \to \Sigma^{-r-s} T^2 E$ of weight $r + s$. Thus, we have a refinement of D with respect

to D' as follows:

$$D'' := \begin{cases} T^{-1}B \to 0 \to B \to B \\ F \to B \to Y \to TF \\ TE \to Y \to X \to \Sigma^{-r-s}TE. \end{cases}$$

Moreover, $w(D'') = r + s = w(D) + w(D')$.

Proof of Proposition 2.55. By definition, the cone decomposition D consists of a family of strict exact triangles in $\mathcal{C}$ as follows:

$$\begin{cases} \vdots \\ \Delta_{i-1} : X_{i-1} \to Y_{i-2} \to Y_{i-1} \to \Sigma^{-r_{i-1}} X_{i-1} \\ \Delta_i : X_i \to Y_{i-1} \to Y_i \to \Sigma^{-r_i} X_i \\ \Delta_{i+1} : X_{i+1} \to Y_i \to Y_{i+1} \to \Sigma^{-r_{i+1}} X_{i+1} \\ \vdots \end{cases}$$

We aim to replace the triangle Δ_i by a sequence of strict exact triangles

$$\bar{\Delta}_j : TA_j \to B_{j-1} \to B_j \to \Sigma^{-s_j} T^2 A_j$$

for $j \in \{1, \dots, k\}$, with $B_0 = Y_{i-1}$, $B_k = Y_i$, and such that

$$\sum_{j=1}^{k} w(\bar{\Delta}_j) = w(\Delta_i) + w(D'). \tag{2.27}$$

In this case, the ordered family of strict exact triangles $(\Delta_1, \dots, \Delta_{i-1}, \bar{\Delta}_1, \dots, \bar{\Delta}_k, \Delta_{i+1}, \dots, \Delta_n)$ forms the refinement D'', and (2.27) implies that

$$w(D'') = \sum_{j \in \{1,\dots,n\} \setminus \{i\}} w(\Delta_j) + \sum_{j=1}^{k} w(\bar{\Delta}_j) = \sum_{j=1}^{n} w(\Delta_j) + w(D') = w(D) + w(D'),$$

as claimed.

In order to obtain the desired sequence of strict exact triangles, we focus on Δ_i and, to shorten notation, we rename its terms by $A = X_i$, $B = Y_{i-1}$, $C = Y_i$, and $r = r_i$ so that, with this notation, Δ_i is a strict exact triangle $A \to B \to C \to \Sigma^{-r} TA$.

We now fix notation for the cone decomposition D' of $A = X_i$. It consists of the following family of strict exact triangles:

$$\begin{cases} \Delta'_1 : A_1 \to 0 \to Z_1 \to \Sigma^{-s_1} TA_1 \\ \Delta'_2 : A_2 \to Z_1 \to Z_2 \to \Sigma^{-s_2} TA_2 \\ \vdots \\ \Delta'_{k-1} : A_{k-1} \to Z_{k-2} \to Z_{k-1} \to \Sigma^{-s_{k-1}} TA_{k-1} \\ \Delta'_k : A_k \to Z_{k-1} \to A \to \Sigma^{-s_k} TA_k. \end{cases}$$

We will apply Proposition 2.49 iteratively. The first step is the following commutative diagram, obtained from (2.23),

$$\begin{array}{ccccc} A_k & \longrightarrow & 0 & \longrightarrow & TA_k \\ \downarrow & & \downarrow & & \downarrow \\ Z_{k-1} & \longrightarrow & B & \longrightarrow & B_{k-1} \\ \downarrow & & \downarrow{\scriptstyle \mathbb{1}_B} & & \downarrow \\ A & \longrightarrow & B & \longrightarrow & C \end{array}$$

for some $B_{k-1} \in \mathrm{Obj}(\mathcal{C})$. Define

$$\overline{\Delta}_k : TA_k \to B_{k-1} \to C \to \Sigma^{-r-s_k} T^2 A_k.$$

We have

$$w(\overline{\Delta}_k) = r + s_k = w(\Delta_i) + w(\Delta'_k). \tag{2.28}$$

We then consider the following commutative diagram, again obtained from (2.23),

$$\begin{array}{ccccc} A_{k-1} & \longrightarrow & 0 & \longrightarrow & TA_{k-1} \\ \downarrow & & \downarrow & & \downarrow \\ Z_{k-2} & \longrightarrow & B & \longrightarrow & B_{k-2} \\ \downarrow & & \downarrow & & \downarrow \\ Z_{k-1} & \longrightarrow & B & \longrightarrow & B_{k-1} \end{array}$$

for some $B_{k-2} \in \mathrm{Obj}(\mathcal{C})$. Define $\overline{\Delta}_{k-1}$ to be the strict exact triangle

$$\overline{\Delta}_{k-1} : TA_{k-1} \to B_{k-2} \to B_{k-1} \to \Sigma^{-s_{k-1}} T^2 A_{k-1}.$$

Then

$$w(\overline{\Delta}_{k-1}) = s_{k-1} = w(\Delta'_{k-1}).$$

Inductively, we obtain $B_i \in \mathrm{Obj}(\mathcal{C})$ and strict exact triangles

$$\overline{\Delta}_i : TA_i \to B_{i-1} \to B_i \to \Sigma^{-s_i} T^2 A_i$$

for $2 \leq i \leq k-1$ such that

$$w(\overline{\Delta}_i) = s_i = w(\Delta'_i). \tag{2.29}$$

The final step lies in considering the diagram

$$\begin{array}{ccccc} A_1 & \longrightarrow & 0 & \longrightarrow & TA_1 \\ \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & B & \longrightarrow & B \\ \downarrow & & \downarrow & & \downarrow \\ Z_1 & \longrightarrow & B & \longrightarrow & B_1 \end{array} \tag{2.30}$$

for some $B \in \mathrm{Obj}(\mathcal{C})$. Define $\overline{\Delta}_1$ to be the strict triangle

$$\overline{\Delta}_1 : TA_1 \to B \to B_1 \to \Sigma^{-s_1} TA_1.$$

Then

$$w(\overline{\Delta}_1) = w(\Delta_1') = s_1. \tag{2.31}$$

Together, the ordered family $(\overline{\Delta}_1, \ldots, \overline{\Delta}_k)$ forms the desired sequence of strict exact triangles. Finally, the equalities (2.28), (2.29), and (2.31) yield

$$\sum_{j=1}^{k} w(\overline{\Delta}_j) = s_1 + s_2 + \cdots + s_k + r = \sum_{j=1}^{k} w(\Delta_j') + w(\Delta_i) = w(D') + w(\Delta_i),$$

as claimed in (2.27). ∎

Let $\mathcal{F} \subset \mathrm{Obj}(\mathcal{C})$ be a family of objects of $\mathcal{C}$. For two objects $X, X' \in \mathrm{Obj}(\mathcal{C})$, define, just as in Section 2.1,

$$\begin{aligned} \delta^{\mathcal{F}}(X, X') = \inf\big\{ w(D) \mid\ & D \text{ is an iterated cone decomposition of } X \\ & \text{with linearization } (F_1, \ldots, T^{-1}X', \ldots, F_k), \\ & \text{where } F_i \in \mathcal{F}, k \in \mathbb{N} \big\}. \end{aligned} \tag{2.32}$$

Corollary 2.57. *With the definition of $\delta^{\mathcal{F}}$ in (2.32), we have the inequality*

$$\delta^{\mathcal{F}}(X, X') \leq \delta^{\mathcal{F}}(X, X'') + \delta^{\mathcal{F}}(X'', X')$$

for any $X, X', X'' \in \mathrm{Obj}(\mathcal{C})$.

Proof. For any $\epsilon > 0$, there are cone decompositions D and D' of X and X'', respectively, such that

$$w(D) \leq \delta^{\mathcal{F}}(X, X'') + \epsilon \quad \text{and} \quad w(D') \leq \delta^{\mathcal{F}}(X'', X') + \epsilon$$

with linearizations

$$\ell(D) = (F_1, \ldots, T^{-1}X'', \ldots, F_s) \quad \text{and} \quad \ell(D') = (F_1', \ldots, T^{-1}X', \ldots, F_k'),$$

respectively, $F_i, F'_j \in \mathcal{F}$. This means that $T^{-1}X''$ has a corresponding cone decomposition $T^{-1}D'$ with linearization $\ell(T^{-1}D') = (T^{-1}F'_1, \dots, T^{-2}X', \dots, T^{-1}F'_k)$. Proposition 2.55 implies that there exists a cone decomposition D'' of X that is a refinement of D with respect to $T^{-1}D'$ such that

$$\ell(D'') = (F_1, \dots, F'_1, \dots, T^{-1}X', \dots, F'_k, \dots, F_s)$$

and

$$w(D'') = w(D) + w(D') \leq \delta^{\mathcal{F}}(X, X'') + \delta^{\mathcal{F}}(X'', X') + 2\epsilon,$$

which implies the claim. ∎

Finally, there are also fragmentation pseudometrics specific to this situation with properties similar to those in Proposition 2.4.

Definition 2.58. Let $\mathcal{C}$ be a triangulated persistence category and let $\mathcal{F} \subset \mathrm{Obj}(\mathcal{C})$. The *fragmentation pseudometric*

$$d^{\mathcal{F}} : \mathrm{Obj}(\mathcal{C}) \times \mathrm{Obj}(\mathcal{C}) \to [0, \infty) \cup \{+\infty\}$$

associated to $\mathcal{F}$ is defined by

$$d^{\mathcal{F}}(X, X') = \max\{\delta^{\mathcal{F}}(X, X'), \delta^{\mathcal{F}}(X', X)\}.$$

Remark 2.59. (a) It is clear from Corollary 2.57 that $d^{\mathcal{F}}$ satisfies the triangle inequality and, by definition, it is symmetric. It is immediate to see that $d^{\mathcal{F}}(X, X) = 0$ for all objects X (this is because of the existence of the exact triangle in $\mathcal{C}_0$, $T^{-1}X \to 0 \to X \to X$). It is of course possible that this pseudometric is degenerate, and it is also possible that it is not finite.

(b) If $X \in \mathcal{F}$, then $\delta^{\mathcal{F}}(TX, 0) = 0$ because of the exact triangle $X \to 0 \to TX \to TX$. On the other hand, $\delta^{\mathcal{F}}(0, TX)$ is not generally trivial. However, the exact triangle $TX \to TX \to 0 \to T^2X$ shows that, if $TX \in \mathcal{F}$, then $\delta^{\mathcal{F}}(0, TX) = 0$.

(c) It follows from the previous point that if $\mathcal{F} = \mathrm{Obj}(\mathcal{C})$, then $d^{\mathcal{F}}(X, X') = 0$ (in other words, the pseudometric $d^{\mathcal{F}}(-, -)$ is completely degenerate). More generally, if the family $\mathcal{F}$ is T invariant (in the sense that if $X \in \mathcal{F}$, then $TX, T^{-1}X \in \mathcal{F}$), then T is an isometry with respect to the pseudometric $d^{\mathcal{F}}$ and $d^{\mathcal{F}}(X, X') = 0$ for all $X, X' \in \mathcal{F}$.

(d) Remark 2.5 (c) applies also in this setting, in the sense that we may define at this triangulated persistence level fragmentation pseudometrics $\underline{d}^{\mathcal{F}}$ given by (the symmetrization of) formula (2.5) but making use of weighted triangles in $\mathcal{C}$ instead of the exact triangles in the triangulated category $\mathcal{D}$.

Recall that by assumption $\mathcal{C}_0$ is triangulated and thus additive. Therefore, for any two objects $X, X' \in \mathrm{Obj}(\mathcal{C}_0) = \mathrm{Obj}(\mathcal{C})$, the direct sum $X \oplus X'$ is a well-defined object in $\mathrm{Obj}(\mathcal{C})$.

Proposition 2.60. *For any $A, B, A', B' \in \mathrm{Obj}(\mathcal{C})$, we have*

$$d^{\mathcal{F}}(A \oplus B, A' \oplus B') \leq d^{\mathcal{F}}(A, A') + d^{\mathcal{F}}(B, B').$$

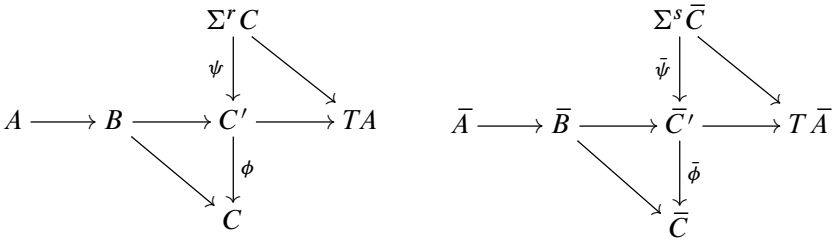
Proof. The proof follows easily from the following lemma.

Lemma 2.61. *Let $\Delta : A \to B \to C \to \Sigma^{-r} TA$ and $\bar{\Delta} : \bar{A} \to \bar{B} \to \bar{C} \to \Sigma^{-s} T\bar{A}$ be two strict exact triangles with $w(\Delta) = r$ and $w(\Delta') = s$. Then*

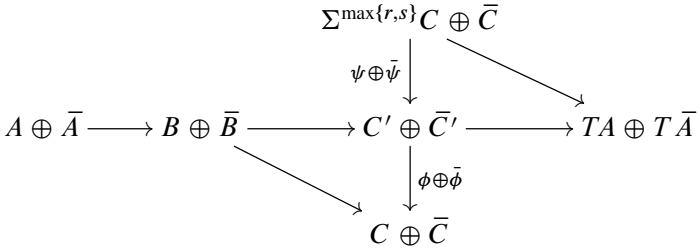
$$\Delta'' : A \oplus \bar{A} \to B \oplus \bar{B} \to C \oplus \bar{C} \to \Sigma^{-\max\{r,s\}}\big(TA \oplus T\bar{A}\big)$$

is a strict exact triangle with $w(\Delta'') = \max\{r, s\}$.

Proof of Lemma 2.61. By definition, there are two commutative diagrams



This yields the commutative diagram



and it is easy to check that $\phi \oplus \bar{\phi}$ is a $\max\{r, s\}$-isomorphism. ∎

Returning to the proof of the proposition, it suffices to prove that $\delta^{\mathcal{F}}(A \oplus B, A' \oplus B') \leq \delta^{\mathcal{F}}(A, A') + \delta^{\mathcal{F}}(B, B')$. For any $\epsilon > 0$, by definition there exist cone decompositions D and D' of A and B, respectively, with $\ell(D) = (F_1, \ldots, F_{i-1}, T^{-1}A', F_{i+1}, \ldots, F_s)$ and $\ell(D') = (F'_1, \ldots, T^{-1}B', \ldots, F'_{s'})$ such that

$$w(D) \leq \delta^{\mathcal{F}}(A, A') + \epsilon \quad \text{and} \quad w(D') \leq \delta^{\mathcal{F}}(B, B') + \epsilon.$$

The desired cone decomposition of $A \oplus B$ is defined as follows:

$$D'' := \begin{cases} F_1 \to 0 \to E_1 \to \Sigma^{-r_1} TF_1 \\ F_2 \to E_1 \to E_2 \to \Sigma^{-r_2} TF_2 \\ \qquad\qquad \vdots \\ F_s \to E_{s-1} \to A \to \Sigma^{-r_s} TF_s \\ F'_1 \to A \oplus 0 \to A \oplus E'_1 \to \Sigma^{-r'_1} TF'_1 \\ F'_2 \to A \oplus E'_1 \to A \oplus E'_2 \to \Sigma^{-r'_2} TF'_{s+2-i} \\ \qquad\qquad \vdots \\ T^{-1}B' \to A \oplus E'_j \to A \oplus E'_{j+1} \to \Sigma^{-r'_{B'}} B' \\ \qquad\qquad \vdots \\ F'_{s'} \to A \oplus E'_{s'-1} \to A \oplus B \to \Sigma^{-r'_{s'}} TF'_{s'}. \end{cases} \tag{2.33}$$

Here we identify $0 \oplus F'_i = F'_i$. The first s triangles come from the decomposition of A, and the following s' triangles are associated, using Lemma 2.61, to the respective triangles in the decomposition of B and to the triangle $0 \to A \to A \to 0$ (of weight 0). It is obvious that $w(D'') = w(D) + w(D')$ and thus $\delta^{\mathcal{F}}(A \oplus B, A' \oplus B') \leq \delta^{\mathcal{F}}(A, A') + \delta^{\mathcal{F}}(B, B')$. ∎

The next statement is an immediate consequence of Proposition 2.60.

Corollary 2.62. *The set* $\mathrm{Obj}(\mathcal{C})$ *with the topology induced by the fragmentation pseudometric* $d^{\mathcal{F}}$ *is an* H*-space relative to the operation*

$$(A, B) \in \mathrm{Obj}(\mathcal{C}) \times \mathrm{Obj}(\mathcal{C}) \to A \oplus B \in \mathrm{Obj}(\mathcal{C}).$$

2.3.4.1 Proof of Proposition 2.4. We now return to the setting in Section 2.1. Thus, $\mathcal{D}$ is a triangulated category, w is a triangular weight on $\mathcal{D}$ in the sense of Definition 2.1, and the quantities $w(D)$ (associated to an iterated cone decomposition D), $\delta^{\mathcal{F}}$, and $d^{\mathcal{F}}$ are defined as in Section 2.1.

The first (and main) step is to establish a result similar to Proposition 2.55. Namely, if X admits an iterated cone decomposition D with linearization $(X_1, \ldots, X_n)$ and some X_i admits a decomposition D' with linearization $(A_1, \ldots, A_k)$, then X admits an iterated cone decomposition D'' with linearization $(X_1, \ldots, X_{i-1}, TA_1, \ldots, TA_k, X_{i+1}, \ldots, X_n)$ and

$$w(D'') \leq w(D) + w(D'). \tag{2.34}$$

For convenience, recall from Section 2.1 that the expression of $w(D)$ is

$$w(D) = \sum_{i=1}^{n} w(\Delta_i) - w_0, \tag{2.35}$$

where $w_0 = w(0 \to X \xrightarrow{\mathbb{1}_X} X \to 0)$ for all X. To show (2.34) we go through exactly the same construction of the refinement D'' of the decomposition D with respect to D', as in the proof of Proposition 2.55, assuming now that all shifts are trivial along the way. The analog of the diagram (2.30) that appears in the last step of the construction of D'' remains possible in this context due to point (ii) of Definition 2.1. By tracking the respective weights along the construction and using Remark 2.2 (b) to estimate the weight of $TA_1 \to B \to B_1 \to T^2 A_1$ from (2.30), we deduce (with the notation in the proof of Proposition 2.55) that

$$\sum_{j=1}^{k} w(\bar{\Delta}_j) \leq w(D') + w(\Delta_i) - w_0,$$

which implies (2.34). Once formula (2.34) is established, it immediately follows that $\delta^{\mathcal{F}}(-,-)$ satisfies the triangle inequality. Further, because the weight of a cone decomposition is given by (2.35), it follows that the cone decomposition D of X with linearization $(T^{-1}X)$ given by the single exact triangle $T^{-1}X \to 0 \to X \to X$ is of weight $w(D) = 0$. As a consequence, $\delta^{\mathcal{F}}(X, X) = 0$. It follows that $d^{\mathcal{F}}(-,-)$ is a pseudometric, as claimed in point (i) of Proposition 2.4.

Assuming now that w is subadditive, the same type of decomposition as in equation (2.33) can be constructed to show that $\delta^{\mathcal{F}}(A \oplus B, A' \oplus B') \leq \delta^{\mathcal{F}}(A, A') + \delta^{\mathcal{F}}(B, B') + w_0$, which implies the claim. ∎

2.4 A persistence triangular weight on $\mathcal{C}_\infty$

The purpose of this section is to further explore the structure of the limit category $\mathcal{C}_\infty$ associated to a triangulated persistence category $\mathcal{C}$. We already know from Section 2.3.2 that the category $\mathcal{C}_\infty$ is triangulated. The main aim here is to use the properties of the weighted exact triangles introduced in Section 2.3.3 to endow $\mathcal{C}_\infty$ with a triangular weight, in the sense of Section 2.1.

2.4.1 Weight of exact triangles in $\mathcal{C}_\infty$

In this section we will use the weighted strict exact triangles in $\mathcal{C}$ to associate weights to the exact triangles in $\mathcal{C}_\infty$.

Assume that $\mathcal{C}$ is a persistence category and recall its ∞-level $\mathcal{C}_\infty$ from Definition 2.9: its objects are the same as those of $\mathcal{C}$ and its morphisms are

$$\hom_{\mathcal{C}_\infty}(A, B) = \varinjlim_{\alpha \to \infty} \hom^\alpha_{\mathcal{C}}(A, B)$$

for any two objects A, B of $\mathcal{C}$. For a morphism $\bar{f}$ in $\mathcal{C}$, we denote by $[\bar{f}]$ the corresponding morphism in $\mathcal{C}_\infty$, and if $f = [\bar{f}]$ for $f \in \hom_{\mathcal{C}_\infty}$ and $\bar{f} \in \hom_{\mathcal{C}}$, we say that $\bar{f}$ represents f. We use the same terminology for diagrams (including triangles) in $\mathcal{C}$ in relation to corresponding diagrams in $\mathcal{C}_\infty$, in the sense that a diagram in $\mathcal{C}$ represents one in $\mathcal{C}_\infty$ if the objects in the two cases are the same and the morphisms in the diagram in $\mathcal{C}$ represent the corresponding ones in the $\mathcal{C}_\infty$ diagram. Clearly, all r-commutativities and r-isomorphisms in $\mathcal{C}$ become, respectively, commutativities and isomorphisms in $\mathcal{C}_\infty$. For instance, if K is r-acyclic, then K is isomorphic to 0 in $\mathcal{C}_\infty$.

For further reference, notice also that the hom-sets of $\mathcal{C}_\infty$ admit a natural filtration as follows. For any $A, B \in \mathrm{Obj}(\mathcal{C}_\infty)$ and $f \in \hom_{\mathcal{C}_\infty}(A, B)$, let the *spectral invariant* of f be given by

$$\sigma(f) := \inf\{k \in \mathbb{R} \cup \{-\infty\} \mid f = [\bar{f}] \text{ for some } \bar{f} \in \hom^k_{\mathcal{C}}(A, B)\} \tag{2.36}$$

and

$$\hom^{\leq \alpha}_{\mathcal{C}_\infty}(A, B) = \{f \in \hom_{\mathcal{C}_\infty}(A, B) \mid \sigma(f) \leq \alpha\}.$$

Assume from now on that $\mathcal{C}$ is a triangulated persistence category. In this case, we have already seen in Section 2.3.2 that $\mathcal{C}_\infty$ is identified with $\mathcal{C}_0/\mathcal{A}\mathcal{C}_0$, the Verdier quotient of $\mathcal{C}_0$ by the subcategory of acyclics. Thus $\mathcal{C}_\infty$ is triangulated with its exact triangles defined through isomorphism with the image in $\mathcal{C}_\infty$ of the exact triangles in $\mathcal{C}_0$; see Remark 2.41.

Before proceeding, we notice that the shift functor Σ associated to $\mathcal{C}$ (see Definition 2.21) induces a similar functor $\Sigma : (\mathbb{R}, +) \to \mathrm{End}(\mathcal{C}_\infty)$. We continue to use the same notation for the r-shifts Σ^r and the natural transformations $\eta_{r,s} : \Sigma^r \to \Sigma^s$. At the same time, in contrast to morphisms in $\mathcal{C}$, there is no meaning to the "amount of shift" for a morphism in $\mathcal{C}_\infty$ (though one can associate to such a morphism its spectral invariant as above). Similarly, the functor T (which is defined as in Remark 2.23 (b) on all of $\mathcal{C}$) also induces a similar functor on $\mathcal{C}_\infty$.

Given a triple of maps $\Delta : A \to B \to C \to D$ in $\mathcal{C}$, the shifted triple $\Sigma^{s_1,s_2,s_3,s_4}\Delta$ was defined in (2.26), and we will use the same notation for similar triples in $\mathcal{C}_\infty$. Note however that in $\mathcal{C}_\infty$ the inequalities relating the s_i's and the shifts of u, v, w are no longer relevant (in fact, do not make sense) and the shift Σ^{s_1,s_2,s_3,s_4} will be used in $\mathcal{C}_\infty$ without these constraints.

Definition 2.63. The *unstable weight* $w_\infty(\Delta)$ of an exact triangle $\Delta : A \xrightarrow{u} B \xrightarrow{v} C \xrightarrow{w} TA$ in $\mathcal{C}_\infty$ is the infimum of the weights of the strict exact triangles in $\mathcal{C}$ of the

form $(\widetilde{\Delta}, r)$, where

$$\widetilde{\Delta} : A \xrightarrow{u'} \Sigma^{-q} B \xrightarrow{v'} \Sigma^{-s} C \xrightarrow{w'} \Sigma^{-r} TA,$$

$0 \leq q \leq s \leq r$, and $\widetilde{\Delta}$ represents Δ in the following sense: the class of the composition $A \xrightarrow{u'} \Sigma^{-q} B \xrightarrow{\eta_{-q,0}} B$ in $\mathcal{C}_\infty$ equals u, and similarly for v' and w'. The *weight* of Δ, $\bar{w}(\Delta)$, is given by

$$\bar{w}(\Delta) = \inf\{w_\infty(\Sigma^{s,0,0,s}\Delta) \mid s \geq 0\}.$$

Remark 2.64. By definition, $\bar{w}(\Delta) \leq w_\infty(\Delta)$, and Example 2.72 below shows that this inequality can be strict.

For the weight $\bar{w}$ of exact triangles in $\mathcal{C}_\infty$ defined as above, recall that $\bar{w}_0$ denotes the normalization constant in Definition 2.1 (ii). The main result is the following.

Theorem 2.65. *Let $\mathcal{C}$ be a triangulated persistence category. The limit category $\mathcal{C}_\infty$ with the triangular structure coming from the identification $\mathcal{C}_\infty \simeq \mathcal{C}_0/\mathcal{AC}_0$ in Proposition 2.39 admits $\bar{w}$ as a triangular, subadditive weight with $\bar{w}_0 = 0$.*

A persistence refinement of a triangulated category $\mathcal{D}$ is a TPC, denoted by $\widetilde{\mathcal{D}}$, such that $\widetilde{\mathcal{D}}_\infty = \mathcal{D}$. The triangular weight $\bar{w}$ as in Theorem 2.65 is called the persistence weight induced by the respective refinement. The following consequence of Theorem 2.65 is immediate from the general constructions in Section 2.1.

Corollary 2.66. *If a small triangulated category $\mathcal{D}$ admits a TPC refinement $\widetilde{\mathcal{D}}$, then* Obj($\mathcal{D}$) *is endowed with a family of fragmentation pseudometrics $\bar{d}^{\mathcal{F}}$, defined as in Section 2.1, associated to the persistence weight $\bar{w}$ induced by the refinement $\widetilde{\mathcal{D}}$, and it has an H-space structure with respect to the topologies induced by these metrics.*

Remark 2.67. (a) We have seen in Section 2.3.4, in particular Definition 2.58, that for a TPC $\mathcal{C}$ there are fragmentation pseudometrics $d^{\mathcal{F}}$ defined on Obj($\mathcal{C}$). The metrics $\bar{d}^{\mathcal{F}}$ associated to the persistence weight $\bar{w}$ on $\mathcal{C}_\infty$ through the construction in Section 2.1 are defined on the same underlying set, Obj($\mathcal{C}$). The relation between them is

$$\bar{d}^{\mathcal{F}} \leq d^{\mathcal{F}}$$

for any family of objects $\mathcal{F}$. The interest in working with $\bar{d}^{\mathcal{F}}$ rather than with $d^{\mathcal{F}}$ is that if $\mathcal{F}$ is a family of triangular generators of $\mathcal{C}_\infty$ and is closed under the action of T, then $\bar{d}^{\mathcal{F}}$ is finite (see Remark 2.5).

(b) As discussed in Remark 2.5 (c) (and also in Remark 2.59 (d)), in this setting too we may define a simpler (but generally larger) fragmentation pseudometric of the type $\underline{\bar{d}}^{\mathcal{F}}$.

We postpone the proof of Theorem 2.65 to Section 2.4.2. Here, we present a few examples shedding some light on Definition 2.63.

Example 2.68. Assume that $\Delta : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r}TA$ is a strict exact triangle of weight r in $\mathcal{C}$. Consider the following triangle in $\mathcal{C}_\infty$,

$$\Delta_\infty : A \xrightarrow{u} B \xrightarrow{v} C \xrightarrow{w} TA,$$

where $u = [\bar{u}]$, $v = [\bar{v}]$, and $w = \left[C \xrightarrow{\bar{w}} \Sigma^{-r}TA \xrightarrow{(\eta_{-r,0})_{TA}} TA\right]$. We claim that Δ_∞ is an exact triangle in $\mathcal{C}_\infty$ and $w_\infty(\Delta_\infty) \leq r$. Indeed, by construction, Δ_∞ is represented by

$$\bar{\Delta}_\infty : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{(\eta_{-r,0})_A \circ \bar{w}} TA$$

and $\lceil \bar{u} \rceil = \lceil \bar{v} \rceil = 0$, $\lceil (\eta_{-r,0})_A \circ \bar{w} \rceil = r + 0 = r$. Now, the shifted triangle $\widetilde{\Delta}_\infty = \Sigma^{0,0,0,-r}\bar{\Delta}_\infty$ obviously equals Δ, the initial strict exact triangle of weight r. Thus, $[\widetilde{\Delta}_\infty] = \Delta_\infty$ and $w(\widetilde{\Delta}_\infty) = r$. Therefore, $w_\infty(\Delta_\infty) \leq r$.

Remark 2.69. A special case of the situation in Example 2.68 is worth emphasizing. Any exact triangle in $\mathcal{C}_0$ induces an exact triangle in $\mathcal{C}_\infty$ with unstable weight equal to 0. This implies that any morphism $f \in \hom_{\mathcal{C}_\infty}(A, B)$ with $\sigma(f) \leq 0$ can be completed to an exact triangle of unstable weight 0 in $\mathcal{C}_\infty$. Indeed, we first represent f by a morphism $\bar{f} \in \hom^{\alpha}_{\mathcal{C}}(A, B)$ with $\alpha \leq 0$. If $\alpha \neq 0$, we shift $\bar{f}$ up using the persistence structure maps and set $\bar{f}' = i_{\alpha,0}(\bar{f})$. We obviously have $[\bar{f}'] = f$. We then complete $\bar{f}'$ to an exact triangle in $\mathcal{C}_0$. The image of this triangle in $\mathcal{C}_\infty$ is exact, of unstable weight 0, and has f as the first morphism in the triple.

Example 2.70. Consider the strict exact triangle $\Delta : A \to 0 \to \Sigma^{-r}TA \xrightarrow{\mathbb{1}} \Sigma^{-r}TA$ in $\mathcal{C}$, which is of weight $r \geq 0$ (see Remark 2.50 (b)). Let Δ_∞ be the following triangle in $\mathcal{C}_\infty$:

$$\Delta_\infty : A \to 0 \to \Sigma^{-r}TA \xrightarrow{\eta_{-r,0}} TA.$$

We claim that if $\sigma(\mathbb{1}_A) = 0$, then $\bar{w}(\Delta_\infty) = w_\infty(\Delta_\infty) = r$. Indeed, assume that $w_\infty(\Delta_\infty) = s < r$. Then there exists a strict exact triangle in $\mathcal{C}$ of the form

$$A \to 0 \to \Sigma^{-r'}TA \xrightarrow{w'} \Sigma^{-s}TA,$$

with $r' \geq r$, of weight s and such that $\lceil w' \rceil = 0$, $[\eta_{-s,0} \circ w' \circ \eta_{-r,-r'}] = [\eta_{-r,0}]$. Notice that $\lceil (\eta_{-s,0} \circ w' \circ \eta_{-r,-r'}) \rceil = r - r' + s$. Thus, as $r' \geq r > s$, we deduce $\sigma([\eta_{-r,0}]) < r$. By writing $\mathbb{1}_A = \eta_{0,-r} \circ \eta_{-r,0}$ we deduce $\sigma(\mathbb{1}_A) < 0$. We conclude that $w_\infty(\Delta_\infty) = r$. We next consider a triangle $\Sigma^{k,0,0,k}\Delta_\infty$:

$$\Sigma^k A \to 0 \to \Sigma^{-r}TA \to \Sigma^k TA$$

and rewrite it as

$$\Delta' : B \to 0 \to \Sigma^{-r-k}TB \to TB,$$

with $B = \Sigma^k A$. Suppose that $w_\infty(\Delta') < r + k$; then $\sigma(\mathbb{1}_B) < 0$, by the previous argument. This implies $\sigma(\mathbb{1}_A) < 0$ and again contradicts our assumption. Thus $w_\infty(\Delta') = r + k$ and $\bar{w}(\Delta_\infty) = r$.

Remark 2.71. For an object $A \in \mathrm{Obj}(\mathcal{C})$ we always have $\sigma(\mathbb{1}_A) \leq 0$. We have seen just above that if $\sigma([(\eta_{-r,0})_A]) < r$, then $\sigma(\mathbb{1}_A) < 0$. The same conclusion remains true if $\sigma([(\eta_{0,-r})_A]) < -r$ by the same argument. Another useful observation is that if the map $\eta_s : \Sigma^s A \to A$ is an r-isomorphism, with $r < s$, then $\sigma(\mathbb{1}_A) \leq r - s$. Indeed, if η_s is an r-isomorphism, then it has a left r-inverse $\psi : \Sigma^r A \to \Sigma^s A$ with $\eta_s \circ \psi = \eta_r$, which implies $\sigma([\eta_r]) \leq \sigma([\eta_s]) \leq -s$ and thus $\sigma(\mathbb{1}_A) \leq r - s$.

Example 2.72. Let $f \in \hom_{\mathcal{C}_\infty}(A, B)$ and suppose that $\sigma(f) > 0$. We will see here that we can extend f to an exact triangle in $\mathcal{C}_\infty$, of unstable weight at most $\sigma(f) + \epsilon$ (for any $\epsilon > 0$) but, at the same time, no triangle extending f has unstable weight less than $\sigma(f)$. Fix $\sigma(f) < t \leq \sigma(f) + \epsilon$ and let $\bar{f} \in \hom^t_{\mathcal{C}}(A, B)$ be a representative of f. Consider the composition $f' : A \xrightarrow{\bar{f}} B \xrightarrow{(\eta_{0,-t})_B} \Sigma^{-t} B$. Then $f' \in \hom^0_{\mathcal{C}}(A, \Sigma^{-t} B)$. There exists an exact triangle in $\mathcal{C}_0$

$$\Delta : A \xrightarrow{f'} \Sigma^{-t} B \xrightarrow{v'} C \xrightarrow{w} TA$$

for some $C \in \mathrm{Obj}(\mathcal{C}_0)$. In particular, $w \in \hom^0_{\mathcal{C}}(C, TA)$. Next, consider the triangle

$$\Sigma^{0,0,0,-t}\Delta : A \xrightarrow{f'} \Sigma^{-t} B \xrightarrow{v'} C \xrightarrow{w'} \Sigma^{-t} TA, \quad \text{where } w' = \eta_t^{TA} \circ w.$$

By Remark 2.50 (b), $\Sigma^{0,0,0,-t}\Delta$ is a strict exact triangle in $\mathcal{C}$ of weight t. Finally, consider the following triangle in $\mathcal{C}$, obtained by shifting up the last three terms of $\Sigma^{0,0,0,-t}\Delta$:

$$\overline{\Delta} : A \xrightarrow{\bar{f} = (\eta_{-t,0})_B \circ f'} B \xrightarrow{v := \Sigma^t v'} \Sigma^t C \xrightarrow{w := \Sigma^t w'} TA.$$

Its image in $\mathcal{C}_\infty$ is the triangle

$$\Delta_\infty : A \xrightarrow{f = [\bar{f}]} B \xrightarrow{[v]} \Sigma^t C \xrightarrow{[w]} TA$$

and is exact. In the terminology of Definition 2.63, the representative $\widetilde{\Delta}$ of Δ_∞ is the strict exact triangle $\Sigma^{0,0,0,-t}\Delta$. In particular, $w_\infty(\Delta_\infty) \leq t$. Notice that Definition 2.63 immediately implies that any triangle $\Delta'' : A \xrightarrow{f} B \to D \to TA$ in $\mathcal{C}_\infty$ satisfies $w_\infty(\Delta'') \geq \sigma(f)$ (because, with the notation of the definition, the weight of the triangle $\widetilde{\Delta}$ in that definition is at least q). At the same time, $\overline{w}(\Delta_\infty) = 0$ because $\Sigma^{t,0,0,t}(\Delta_\infty)$ has as representative

$$\Sigma^t \Delta : \Sigma^t A \to B \to \Sigma^t C \to \Sigma^t TA,$$

which is exact in $\mathcal{C}_0$.

Remark 2.73. (a) The definition of the weights of the exact triangles in $\mathcal{C}_\infty$ is designed precisely to allow for the construction in Example 2.72. This is quite different compared to the case when the spectral invariant of f is non-positive (compare

with Remark 2.69) because the persistence structure maps can be used to "shift" up but not down. It also follows from Example 2.72 that for $f \in \hom_{\mathcal{C}_\infty}(A, B)$ with $\sigma(f) \geq 0$ we have

$$\sigma(f) = \inf\{w_\infty(\Delta) \mid \Delta : A \xrightarrow{f} B \to C \to TA\}.$$

(b) It is not difficult to see that if $w_\infty(\Sigma^{s,0,0,s}\Delta) \leq r$ for a triangle Δ in $\mathcal{C}_\infty$ and $s \geq 0$, then $w_\infty(\Delta) \leq r + s$.

Example 2.74. Let $\Delta : A \to B \to C \to TA$ be an exact triangle in $\mathcal{C}_0$. It is clear that, in $\mathcal{C}_\infty$ and for $s \leq 0$, the corresponding triangles of the form $\Delta_s : \Sigma^s A \to B \to C \to \Sigma^s TA$ (defined using the (pre)-composition of the maps in Δ with the appropriate maps η's on A) have the property $\overline{w}(\Delta_s) = 0$. On the other hand, if $s > 0$ this is no longer the case, in general. Indeed, assuming $s > 0$ and $\overline{w}(\Delta_s) = 0$, it follows that for some possibly even larger $r > 0$ we have $w_\infty(\Delta_r) = 0$. This means that for any sufficiently small $\epsilon > 0$ there is an exact triangle in $\mathcal{C}_0$ of the form $\Sigma^{r'}A \to \Sigma^{-\epsilon'}B \to C' \to \Sigma^{r'}TA$, where $r' \geq r$, together with an ϵ-isomorphism $\phi : C' \to \Sigma^{-\epsilon''}C$ satisfying the conditions in Definition 2.42, with $0 \leq \epsilon' \leq \epsilon'' \leq \epsilon$ (see Definition 2.63). This triangle can be compared to a shift $\Sigma^{-2\epsilon}\Delta$ of the initial Δ (the constant 2 is necessary to ensure the commutativity of the rectangles in the comparison diagram; see Remark 2.43 (c)). The resulting commutative diagram is

$$\begin{array}{ccccccc}
\Sigma^{r'}A & \longrightarrow & \Sigma^{-\epsilon'}B' & \longrightarrow & C' & \longrightarrow & \Sigma^{r'}TA \\
\downarrow{\scriptstyle \eta^A} & & \downarrow{\scriptstyle \eta^B} & & \downarrow{\scriptstyle \eta^C_{2\epsilon}\circ\phi} & & \downarrow{\scriptstyle \eta^A} \\
\Sigma^{-2\epsilon}A & \longrightarrow & \Sigma^{-2\epsilon}B & \longrightarrow & \Sigma^{-2\epsilon}C & \longrightarrow & \Sigma^{-2\epsilon}TA
\end{array}$$

The map out of C' in this diagram is a 3ϵ-isomorphism and, by Proposition 2.31 (ii), we deduce that $\eta^A_{r'+2\epsilon}$ is a 15ϵ-isomorphism. We can take ϵ as small as needed and we deduce $\sigma(\mathbb{1}_A) \leq -r' \leq -s$.

Example 2.75. Let $\Delta : A \to 0 \to \Sigma^r TA \to TA$ be a triangle in $\mathcal{C}_\infty$ with the last map (the class of) η^{TA}_r and with $r \geq 0$. Assuming $\sigma(\mathbb{1}_A) = 0$, we claim that $w_\infty(\Delta) = r$ and $\overline{w}(\Delta) = 0$. Indeed, the fact that $\overline{w}(\Delta) = 0$ is obvious because $\Sigma^{r,0,0,r}\Delta : \Sigma^r A \to 0 \to \Sigma^r TA \to \Sigma^r TA$ is exact in $\mathcal{C}_0$. Now assume that $w_\infty(\Delta) < r$. Then there is a triangle $A \to 0 \to \Sigma^{r-s'}TA \xrightarrow{u} \Sigma^{-s}TA$ which is strict exact in $\mathcal{C}$ and with $s \geq s'$ and $w_\infty(\Delta) \leq s < r$. Using the definition of strict exact triangles and the existence of the exact triangle $A \to 0 \to TA \to TA$ in $\mathcal{C}_0$, we deduce that there exists an s-isomorphism $\phi : TA \to \Sigma^{r-s'}TA$ with a right inverse $\psi : \Sigma^{s+r-s'}TA \to TA$ that coincides with η^{TA} in $\mathcal{C}_\infty$. As a result, ϕ coincides with $\eta^{TA}_{0,r-s'}$ in $\mathcal{C}_\infty$ and thus $0 = \lceil\phi\rceil \geq r - s'$ ($\sigma(\mathbb{1}_A) = 0$ implies that $\sigma\left([\eta^{TA}_{0,r-s'}]\right) = r - s'$). Therefore, $s' \geq r$, which contradicts our assumption $s < r$.

Example 2.76. Another example of interest is given by the triangle in $\mathcal{C}_\infty$:

$$\Delta : 0 \to X \to \Sigma^r X \to 0,$$

where the map $X \to \Sigma^r X$ is the class of $(\eta_{0,r})_X$. We claim that $w_\infty(\Delta) \leq |r|$. We start with the case $r > 0$. In that case, by shifting down the last two objects by r, we obtain a strict exact triangle in $\mathcal{C}$: $(0 \to X \to X \to 0, r)$ (here it is important to view strict exact triangles as pairs (triangle, weight) as in Definition 2.42). In the case $r \leq 0$, there is a strict exact triangle in $\mathcal{C}$: $(0 \to X \to \Sigma^r X \to 0, -r)$ that can be reached from Δ by shifting down the last object by $|r|$. This shows the claim. Additionally, it is easy to see that if $\sigma(\mathbb{1}_X) = 0$, then $w_\infty(\Delta) = r$.

Example 2.77. In this example we consider an exact triangle in $\mathcal{C}_\infty$,

$$\Delta : 0 \to X \to Y \to 0.$$

We claim that $\overline{w}(\Delta) = \inf\{r \geq 0 \mid \exists \text{ an } r\text{-isomorphism } \phi : \Sigma^s X \to Y \text{ with } 0 \leq s \leq r\}$. By the definition of the weight of triangles in $\mathcal{C}_\infty$, we are looking for strict exact triangles in $\mathcal{C}$ of the form

$$(0 \to X \to \Sigma^{-s} Y \to \Sigma^{-r} 0, r)$$

with $s \leq r$. The unstable weight, in this case equal to the stable weight, is obtained by infimizing r. The existence of such a strict exact triangle is equivalent to the existence of an r-isomorphism $\phi : X \to \Sigma^{-r} Y$, which shows the claim.

Example 2.78. Consider an exact triangle in $\mathcal{C}_\infty$ of the form

$$\Delta : T^{-1} X \to 0 \to Y \to X.$$

We claim that $w_\infty(\Delta) = \inf\{r \geq 0 \mid \exists \text{ an } r\text{-isomorphism } \phi : \Sigma^s X \to Y \text{ with } 0 \leq s \leq r\}$ and $\overline{w}(\Delta) = \inf\{r \geq 0 \mid \exists \text{ an } r\text{-isomorphism } \phi : \Sigma^s X \to Y \text{ with } s \geq 0\}$. The relevant strict exact triangles (of weight r) in this case are

$$T^{-1} X \to 0 \to \Sigma^{-s} Y \to \Sigma^{-r} X$$

with $s \leq r$. Again, the existence of such a triangle is equivalent to the existence of an r-isomorphism $\phi : X \to \Sigma^{-s} Y$, which provides the estimate for w_∞. To estimate $\overline{w}$, we apply the same argument but replace X by $\Sigma^k X$ with k positive. The condition then becomes that $\phi : \Sigma^k X \to \Sigma^{-s} Y$ is an r-isomorphism and that $s \leq r$.

Example 2.79. Consider an exact triangle in $\mathcal{C}_\infty$ of the form

$$\Delta : T^{-1} X \to K \to Y \to X,$$

where K is k-acyclic. We claim that

$$\overline{w}(\Delta) \geq \inf\{r \geq 0 \mid \exists \text{ an } (r+2k)\text{-isomorphism } \phi : \Sigma^s X \to Y \text{ with } s \geq 0\}.$$

We thus assume that there exists a strict exact triangle $\Delta' : T^{-1}X \to \Sigma^{-s_1}K \to \Sigma^{-s_2}Y \to \Sigma^{-r}X$ of weight r, with $s_1 \leq s_2 \leq r$, that represents Δ. As this is strict exact, there is an exact triangle in $\mathcal{C}_0$ of the form $T^{-1}X \to \Sigma^{-s_1}K \to C \to X$ that maps to Δ as in Definition 2.42. In particular, there is an r-isomorphism $C \to \Sigma^{-s_2}Y$. The fact that $K \simeq_k 0$ immediately implies that there is a k-isomorphism $C \to X$. We consider a right k-inverse of it, $\Sigma^k X \to C$. This a $2k$-isomorphism. By composition we get a $(2k+r)$-isomorphism $\Sigma^k X \to \Sigma^{-s_2}Y$. The constraint $s_2 \leq r$ is eliminated by applying the same argument as in Example 2.78, replacing X with $\Sigma^s X$ for some non-negative s.

Remark 2.80. It is useful to know how the weights of triangles in $\mathcal{C}_\infty$ behave with respect to rotation. Thus let $\Delta : A \xrightarrow{u} B \xrightarrow{v} C \xrightarrow{w} TA$ be an exact triangle in $\mathcal{C}_\infty$ of unstable weight r. Then its first positive rotation $R(\Delta) : B \xrightarrow{v} C \xrightarrow{w} TA \xrightarrow{-Tu} TB$ is of unstable weight at most $2r$. Indeed, by definition, there exists a triangle $\overline{\Delta} : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} TA$ representing Δ such that the shifted triangle $\widetilde{\Delta} : A \xrightarrow{u'} \Sigma^{-t_1}B \xrightarrow{v'} \Sigma^{-t_1-t_2}C \xrightarrow{w'} \Sigma^{-r}TA$ (where we put $t_1 = \lceil \bar{u} \rceil$, $t_2 = \lceil \bar{v} \rceil$, $t_3 = \lceil \bar{w} \rceil$, and $r = t_1 + t_2 + t_3$) is a strict exact triangle in $\mathcal{C}$ of weight r. By Proposition 2.46, the first positive rotation

$$R(\widetilde{\Delta}) : \Sigma^{-t_1}B \xrightarrow{v'} \Sigma^{-t_1-t_2}C \xrightarrow{w''} \Sigma^{-r}TA \xrightarrow{u''} \Sigma^{-t_1-2r}TB$$

is a strict exact triangle in $\mathcal{C}$ of weight $2r$. We have

$$\begin{aligned} u'' &= -(\eta_{0,-t_1-2r})_B \circ \bar{u} \circ (\eta_{-r,0})_A \\ v' &= (\eta_{0,-t_1-t_2})_C \circ \bar{v} \circ (\eta_{-t_1,0})_B \\ w'' &\simeq_r w = (\eta_{0,-r})_A \circ \bar{w} \circ (\eta_{-t_1-t_2,0})_C. \end{aligned}$$

Consider a new triangle:

$$\Sigma^{t_1}R(\widetilde{\Delta}) : B \xrightarrow{\Sigma^{t_1}v'} \Sigma^{-t_2}C \xrightarrow{\Sigma^{t_1}w''} \Sigma^{-r+t_1}TA \xrightarrow{\Sigma^{t_1}u''} \Sigma^{-2r}TB.$$

By Remark 2.50 (a), $\Sigma^{t_1}R(\widetilde{\Delta})$ is also a strict exact triangle in $\mathcal{C}$ of weight $2r$. By shifting this triangle up we get to $B \to C \to TA \to TB$, which represents $R(\Delta)$; thus $w_\infty(R(\Delta)) \leq 2r$. In a similar way, using Remark 2.47, one can treat the negative rotation $R^{-1}(\Delta)$ of Δ.

2.4.2 Proof of Theorem 2.65

There are two steps. The first is to show that each exact triangle in $\mathcal{C}_\infty$ has finite unstable weight. The second step is to show that $\overline{w}$ satisfies Definition 2.1 and is subadditive with $\overline{w}_0 = 0$.

2.4.2.1 Every triangle in $\mathcal{C}_\infty$ has finite unstable weight. Let $\Delta : A \xrightarrow{u} B \xrightarrow{v} C \xrightarrow{w} TA$ be an exact triangle in $\mathcal{C}_\infty$. Thus, there exist a triangle $\overline{\Delta} : A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} TA$ in $\mathcal{C}$ and an exact triangle in $\mathcal{C}_0$, $\Delta' : A' \xrightarrow{u'} B' \xrightarrow{v'} C' \xrightarrow{w'} TA'$, together with isomorphisms in $\mathcal{C}_\infty$, $a : A' \to A$, $b : B' \to B$, $c : C' \to C$, such that the resulting map $(a, b, c, -Ta)$ of triangles in $\mathcal{C}_\infty$ is an isomorphism of triangles. We may assume that the shifts x of $\bar{u}$, y of $\bar{v}$, and z of $\overline{w}$ are all non-negative.

We represent the maps a, b, c by s-isomorphisms $\bar{a} : \Sigma^r A' \to A$, $\bar{b} : \Sigma^r B' \to B$, $\bar{c} : \Sigma^r C' \to C$, which is possible by taking s and r sufficiently large. We now consider the diagram in $\mathcal{C}_0$:

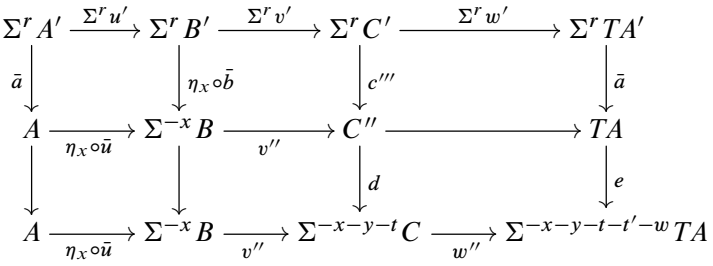
$$\begin{array}{ccc} \Sigma^r A' & \xrightarrow{\Sigma^r u'} & \Sigma^r B' \\ \downarrow{\scriptstyle \bar{a}} & & \downarrow{\scriptstyle \eta_x \circ \bar{b}} \\ A & \xrightarrow[\eta_x \circ \bar{u}]{} & \Sigma^{-x} B \end{array}$$

This diagram r'-commutes for r' sufficiently large. Lemma 2.17 shows that, by taking r sufficiently large, we may assume that the above square commutes in $\mathcal{C}_0$ (at the cost of making both r and s very large). We complete the diagram to the right, thus obtaining a new diagram in $\mathcal{C}_0$:

$$\begin{array}{ccccccc} \Sigma^r A' & \xrightarrow{\Sigma^r u'} & \Sigma^r B' & \xrightarrow{\Sigma^r v'} & \Sigma^r C' & \xrightarrow{\Sigma^r w'} & \Sigma^r TA' \\ \downarrow{\scriptstyle \bar{a}} & & \downarrow{\scriptstyle \eta_x \circ \bar{b}} & & \downarrow{\scriptstyle c''} & & \downarrow{\scriptstyle a''} \\ A & \xrightarrow[\eta_x \circ \bar{u}]{} & \Sigma^{-x} B & \xrightarrow[v'']{} & \Sigma^{-x-y-t} C & \xrightarrow[w'']{} & \Sigma^{-x-y-t-t'-w} TA \end{array}$$

with each square commutative and with the arrows marked with $(-)''$ being compositions of $(-)$ with the appropriate η's. The commutativity in $\mathcal{C}_0$ for the second square requires possibly a shift by $-t$ of the third term in the bottom row. Similarly, the commutativity of the third square requires a shift by $-t'$ for the last term of the row (this is a variant of Proposition 2.51). Obviously, the top row is exact in $\mathcal{C}_0$ and the vertical maps are k-isomorphisms, each for a different value of k. At the cost of yet again increasing t and t', we can intercalate between the two rows a new exact

triangle in $\mathcal{C}_0$:

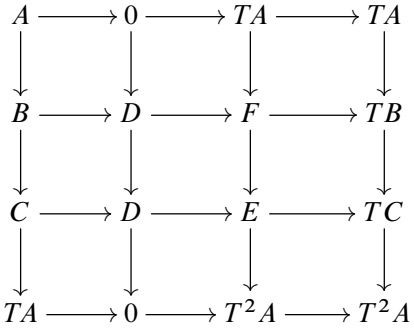


The map c''' is induced by the first two vertical maps between the top rows, and thus it is a k-isomorphism for some sufficiently large k. Given that c'' is such an isomorphism too, by possibly increasing t we obtain the existence of a k'-isomorphism d with k' very large. By possibly increasing t' we can also get the commutativity of the bottom-right square. Again by possibly increasing t' we may ensure that $k' = x + y + t + t' + w$ (recall that any k'-isomorphism is also a k''-isomorphism for $k'' \geq k'$). This means that the bottom row is a strict exact triangle in $\mathcal{C}$ of weight k'. To end the proof, we notice that this triangle is of the form $\widetilde{\Delta}$ as in Definition 2.63.

2.4.2.2 The weighted octahedral axiom in $\mathcal{C}_\infty$. To finish the proof of Theorem 2.65 we need to show that the weighted octahedral axiom is satisfied by $\overline{w}$, that $\overline{w}$ satisfies the normalization in Definition 2.1 (ii) with $\overline{w}_0 = 0$, and that it is subadditive. We start below with the weighted octahedral axiom and will end with the other properties.

Lemma 2.81. *The weight $\overline{w}$ as defined in Definition 2.63 satisfies the weighted octahedral axiom from Definition 2.1* (i).

Proof. Recall that given the exact triangles $\Delta_1 : A \to B \to C \to TA$ and $\Delta_2 : C \to D \to E \to TC$ in $\mathcal{C}_\infty$ we need to show that there are exact triangles $\Delta_3 : B \to D \to F \to TB$ and $\Delta_4 : TA \to F \to E \to T^2A$ making the diagram below commute, except for the bottom-right square, which anti-commutes:



and such that $\overline{w}(\Delta_3) + \overline{w}(\Delta_4) \leq \overline{w}(\Delta_1) + \overline{w}(\Delta_2)$.

In $\mathcal{C}$, there are triangles $\overline{\Delta}_1 : A \to B \to C \to TE$ with non-negative morphism shifts t_1, t_2, t_3, and $\overline{\Delta}_2 : C \to D \to E \to TC$ with non-negative morphism shifts k_1, k_2, k_3, representing Δ_1 and Δ_2, respectively, and such that the associated triangles

$$\widetilde{\Delta}_1 : A \to \Sigma^{-t_1} B \to \Sigma^{-t_1-t_2} C \to \Sigma^{-r} TA, \quad \text{where } r = t_1 + t_2 + t_3,$$

and

$$\widetilde{\Delta}_2 : C \to \Sigma^{-k_1} D \to \Sigma^{-k_1-k_2} E \to \Sigma^{-s} TC, \quad \text{where } s = k_1 + k_2 + k_3,$$

are strict exact triangles in $\mathcal{C}$ of weight r and s, respectively. Consider $\Sigma^{-t_1-t_2}\widetilde{\Delta}_2$:

$$\Sigma^{-t_1-t_2}\widetilde{\Delta}_2 : \Sigma^{-t_1-t_2} C \to \Sigma^{-k_1-t_1-t_2} D \to \Sigma^{-k_1-k_2-t_1-t_2} E \to \Sigma^{-s-t_1-t_2} TC.$$

The weighted octahedral property for strict exact triangles in $\mathcal{C}$, stated in Proposition 2.49, implies that we can construct the following diagram in $\mathcal{C}$, commutative except for the bottom-right square, which is r-anti-commutative:

$$\begin{array}{ccccccc}
A & \longrightarrow & 0 & \longrightarrow & TA & \longrightarrow & TA \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\Sigma^{-t_1} B & \longrightarrow & \Sigma^{-k_1-t_1-t_2} D & \longrightarrow & D' & \longrightarrow & \Sigma^{-t_1} TB \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\Sigma^{-t_1-t_2} C & \longrightarrow & \Sigma^{-k_1-t_1-t_2} D & \longrightarrow & \Sigma^{-k_1-k_2-t_1-t_2} E & \longrightarrow & \Sigma^{-t_1-t_2-s} TC \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\Sigma^{-r} TA & \longrightarrow & 0 & \longrightarrow & \Sigma^{-r-s} T^2 A & \longrightarrow & \Sigma^{-r-s} T^2 A
\end{array} \tag{2.37}$$

Here the triangle $\Delta'_3 : \Sigma^{-t_1} B \to \Sigma^{-k_1-t_1-t_2} D \to D' \to \Sigma^{-t_1} TB$ is an exact triangle in $\mathcal{C}_0$. The triangle $\Delta''_3 : \Sigma^{k_1+t_2} B \to D \to \Sigma^{k_1+t_1+t_2} D' \to \Sigma^{k_1+t_2} TB$, obtained by shifting Δ'_3 up by $k_1 + t_1 + t_2$, is also exact in $\mathcal{C}_0$. Let $[\Delta''_3]$ be the image of this triangle in $\mathcal{C}_\infty$. We put $F = \Sigma^{k_1+t_1+t_2} D'$ and take Δ_3 to be the triangle in $\mathcal{C}_\infty$

$$\Delta_3 : B \to D \to F \to TB$$

obtained by applying $\Sigma^{-k_1-t_2,0,0,-k_1-t_2}$ to $[\Delta''_3]$. We obviously have $w_\infty([\Delta''_3]) = 0$ and thus $\overline{w}(\Delta_3) = 0$.

The next step is to identify the triangle Δ_4. Proposition 2.49 implies that the third column in (2.37),

$$\Delta'_4 : TA \to \Sigma^{-k_1-t_1-t_2} F \to \Sigma^{-k_1-k_2-t_1-t_2} E \to \Sigma^{-r-s} T^2 A,$$

is a strict exact triangle in $\mathcal{C}$. Let $[\Delta'_4]$ be the image of this triangle in $\mathcal{C}_\infty$, and let Δ_4 be given by applying $\Sigma^{0,k_1+t_1+t_2,k_1+k_2+t_1+t_2,r+s}$ to $[\Delta'_4]$:

$$\Delta_4 : TA \to F \to E \to T^2 A.$$

We deduce from Definition 2.63 that $w_\infty(\Delta_4) \leq r + s$ and thus also $\overline{w}(\Delta_4) \leq r + s$.

The commutativity required in the statement follows from that provided by Proposition 2.49 for (2.37). ∎

Remark 2.82. For the triangle Δ_3 produced in the proof above, it is easy to see that $w_\infty(\Delta_3) \leq k_1 + t_2$. Therefore we have

$$w_\infty(\Delta_3) + w_\infty(\Delta_4) \leq (k_1 + t_2) + (r + s) \leq 2(r + s) = 2(w_\infty(\Delta_1) + w_\infty(\Delta_2)).$$

Thus the weight w_∞ satisfies a weak form of the weighted octahedral axiom.

The next step in proving Theorem 2.65 is to show the normalization property in Definition 2.1 (ii). This property is satisfied with the constant $\overline{w}_0 = 0$. Indeed, any triangle $0 \to X \to X \to 0$ and all its rotations are exact in $\mathcal{C}_0$, and thus they are of unstable weight equal to 0. The last verification needed is to see that, if $B = 0$ in the diagram of the weighted octahedral axiom, then the triangle Δ_3 (constructed in the proof of Lemma 2.81) can be of the form $\Delta_3 : 0 \to D \to D \to 0$. This is trivially satisfied in our construction because if $B = 0$ we may take $t_1 = t_2 = 0$ and the triangle $\Delta_3' : 0 \to \Sigma^{-k_1} D \xrightarrow{\mathbb{1}} \Sigma^{-k_1} D \to 0$.

Finally, to finish the proof of Theorem 2.65 we need to show that $\overline{w}$ is subadditive. Thus, assuming that $\Delta : A \to B \to C \to TA$ is exact in $\mathcal{C}_\infty$ and that X is an object in $\mathcal{C}$, we have $\overline{w}(X \oplus \Delta) \leq \overline{w}(\Delta)$, where the triangle $X \oplus \Delta$ has the form $X \oplus \Delta : A \to X \oplus B \to X \oplus C \to TA$. We consider the strict exact triangle in $\mathcal{C}$

$$\widetilde{\Delta} : A \to \Sigma^{-s_1} B \to \Sigma^{-s_1 - s_2} C \to \Sigma^{-s_1 - s_2 - s_3} TA$$

associated to Δ as in Definition 2.63, where $s_i \geq 0$, $1 \leq i \leq 3$. Consider the triangle

$$\Delta' : 0 \to \Sigma^{-s_1} X \xrightarrow{\eta^X_{s_2}} \Sigma^{-s_1 - s_2} X \to 0.$$

This triangle is obtained from the exact triangle in $\mathcal{C}_0$ $0 \to \Sigma^{-s_1} X \xrightarrow{\mathbb{1}} \Sigma^{-s_1} X \to 0$ by applying $\Sigma^{0,0,-s_2,-s_2}$, and it is of weight at most s_2. By Lemma 2.61 we have $w(\Delta' \oplus \widetilde{\Delta}) \leq w(\widetilde{\Delta})$. We now note that $\Delta' \oplus \widetilde{\Delta}$ can be viewed as being obtained from $X \oplus \Delta$ by applying $\Sigma^{0,-s_1,-s_1-s_2,-s_1-s_2-s_3}$, and thus $w_\infty(X \oplus \Delta) \leq w(\widetilde{\Delta})$, which implies the claim. The proof for $\Delta \oplus X$ is similar.

2.4.3 Some properties of fragmentation pseudometrics

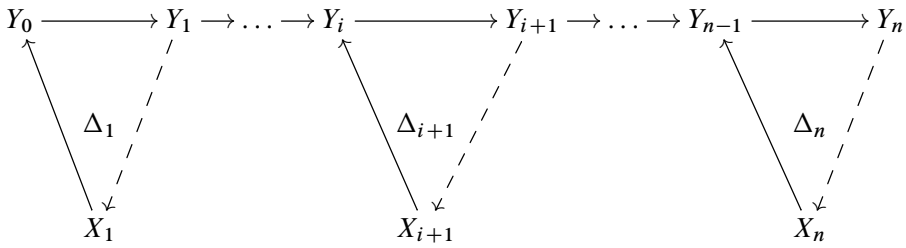
The purpose of this section is to rapidly review some of the properties of the persistence fragmentation pseudometrics. We start by recalling the main definitions; we then discuss some algebraic properties and relations to standard notions such as the interleaving distance.

2.4.3.1 Persistence fragmentation pseudometrics: Review of main definitions. Assume that $\mathcal{C}$ is a TPC. We have defined in Sections 2.3.4 and 2.4.1 three similarly defined types of measurements on the objects of $\mathcal{C}$ which, after symmetrization, define fragmentation pseudometrics on $\mathrm{Obj}(\mathcal{C})$. In Section 2.3.4, this construction uses directly the weight w of the strict exact triangles in $\mathcal{C}$ (making use of Proposition 2.49), and it gives rise to the pseudometrics $d^{\mathcal{F}}$ as in Definition 2.58, as well as to a simplified version $\underline{d}^{\mathcal{F}}$ mentioned in Remark 2.59 (d).

In Section 2.4.1 we have endowed the triangles of the category $\mathcal{C}_\infty$ with weights: an unstable weight w_∞ as well as a (smaller) stable weight $\overline{w}$, as given in Definition 2.63. Working in the category $\mathcal{C}_\infty$ has a significant advantage over working in the category $\mathcal{C}$ because, in contrast to the latter, in $\mathcal{C}_\infty$ any morphism can be completed to an exact triangle of finite weight. Moreover, when we have a (translation-invariant) family of triangular generators $\mathcal{F}$, the fragmentation pseudodistance relative to $\mathcal{F}$ is finite; see Remark 2.5 (a). Finally, in $\mathcal{C}_\infty$ exact triangles have the standard form expected in a triangulated category and do not involve shifts. Therefore, we will focus here on the fragmentation pseudometrics defined using w_∞ and, mainly, $\overline{w}$.

An important remark at this point is that the unstable weight w_∞ does not satisfy the weighted octahedral axiom (but only its weak form, as discussed in Remark 2.82), and thus only the pseudometrics of the form $\underline{\bar{d}}^{\mathcal{F}}$ can be defined using it. By contrast, $\overline{w}$ does satisfy the weighted octahedral axiom and there is a pseudometric $\bar{d}^{\mathcal{F}}$ associated to it through the construction in Section 2.1. Both $\overline{w}$ and w_∞ are subadditive and satisfy the normalization property in Definition 2.1 with $w_0 = 0$.

To eliminate possible ambiguities, we recall the definitions of the two relevant pseudometrics here. Both of them are based on considering a sequence of exact triangles in $\mathcal{C}_\infty$ as below:



(2.38)

where the dotted arrows represent maps $Y_i \to TX_i$. We fix a family of objects $\mathcal{F}$ in $\mathcal{C}$ with $0 \in \mathcal{F}$ and define

$$\underline{\bar{\delta}}^{\mathcal{F}}(X, X') = \inf\Big\{\sum_{i=1}^{n} w_\infty(\Delta_i) \mid \Delta_i \text{ are successive exact triangles as in (2.38)} \\ \text{with } X' = Y_0,\ X = Y_n, \text{ and } X_i \in \mathcal{F}, n \in \mathbb{N}\Big\}, \tag{2.39}$$

and

$$\bar{\delta}^{\mathcal{F}}(X, X') = \inf\Big\{\sum_{i=1}^{n} \overline{w}(\Delta_i) \mid \Delta_i \text{ are as in (2.38) with } Y_0 = 0,\ X = Y_n,\ X_i \in \mathcal{F},\ n \in \mathbb{N} \text{ except for some } j \text{ such that } X_j = T^{-1}X'\Big\}. \tag{2.40}$$

Finally, the pseudometrics $\underline{\bar{d}}^{\mathcal{F}}$ and $\bar{d}^{\mathcal{F}}$ are obtained by symmetrizing $\underline{\bar{\delta}}^{\mathcal{F}}$ and $\bar{\delta}^{\mathcal{F}}$, respectively:

$$\underline{\bar{d}}^{\mathcal{F}}(X, X') = \max\{\underline{\bar{\delta}}^{\mathcal{F}}(X, X'), \underline{\bar{\delta}}^{\mathcal{F}}(X', X)\},$$
$$\bar{d}^{\mathcal{F}}(X, X') = \max\{\bar{\delta}^{\mathcal{F}}(X, X'), \bar{\delta}^{\mathcal{F}}(X', X)\}.$$

2.4.3.2 Algebraic properties. There are many fragmentation pseudometrics of persistence type associated to the same weight, depending on the choice of family $\mathcal{F}$. In fact, the choices available are even more abundant for the following two reasons:

(i) Triangular weights themselves can be mixed. For instance, if $\mathcal{C}$ is a TPC, there is a triangular weight of the form $\overline{w}^{+} = \overline{w} + w_{fl}$ that is defined on $\mathcal{C}_\infty$ (where w_{fl} is the flat weight defined in Section 2.1).

(ii) Fragmentation metrics themselves can also be mixed. If $d^{\mathcal{F}}$ and $d^{\mathcal{F}'}$ are two fragmentation pseudometrics (whether defined with respect to the same weight or not), then the expressions $\alpha d^{\mathcal{F}} + \beta d^{\mathcal{F}'}$ with $\alpha, \beta \geq 0$, as well as $\max\{d^{\mathcal{F}}, d^{\mathcal{F}'}\}$, are also pseudometrics.

In essence, although it is not easy to produce interesting subadditive triangular weights on a triangulated category, once such a weight is constructed – as in the case of the persistence weight $\overline{w}$ defined on $\mathcal{C}_\infty$ (where $\mathcal{C}$ is a TPC) – one can associate to it a large class of pseudometrics, either by combining the weight with the flat one and/or by "mixing" the pseudometrics associated to different families $\mathcal{F}$. Another useful (and obvious) property relating the pseudometrics $d^{\mathcal{F}}$ and $d^{\mathcal{F}'}$ associated to the same triangular weight is the following:

(iii) If $\mathcal{F} \subset \mathcal{F}'$, then $d^{\mathcal{F}'} \leq d^{\mathcal{F}}$.

The last useful construction has to do with making the metrics invariant with respect to the action of the shift functor.

(iv) For a given fragmentation metric $d^{\mathcal{F}}$ we define its shift-invariant version

$$\widehat{d^{\mathcal{F}}}(X, Y) = \max\{\widehat{\delta^{\mathcal{F}}}(X, Y), \widehat{\delta^{\mathcal{F}}}(Y, X)\}. \tag{2.41}$$

Here

$$\widehat{\delta^{\mathcal{F}}}(X, Y) = \inf_{r,s \in \mathbb{R}} \delta^{\mathcal{F}}(\Sigma^r X, \Sigma^s Y)$$

is the shift-invariant version of the semi (pseudo)metrics $\delta^{\mathcal{F}}$ as in (2.39) and (2.40). It is immediate to see that $\widehat{\delta^{\mathcal{F}}}$ satisfies the triangle inequality. By symmetrizing, we

obtain indeed a pseudometric that is obviously bounded above by $d^{\mathcal{F}}$. In case the family $\mathcal{F}$ is closed under the action of Σ^r for all $r \in \mathbb{R}$, the metrics of type $d^{\mathcal{F}}$ have the property that $d^{\mathcal{F}}(X, Y) = d^{\mathcal{F}}(\Sigma^r X, \Sigma^r Y)$ for all $r \in \mathbb{R}$. In this case, the shift-invariant metric associated to $d^{\mathcal{F}}$ has a simpler form $\widehat{d^{\mathcal{F}}}(X, Y) = \inf_{r\in\mathbb{R}} d^{\mathcal{F}}(\Sigma^r X, Y)$. The interest of this type of shift-invariant pseudometric is that it compares the "shape" of objects, in contrast to a comparison of the objects themselves that is sensitive to translations (the spectral distance in symplectic topology is of this type). Thus, for instance, two Morse functions f and $f + k$ with $k \in \mathbb{R}$ are not distinguished by shift-invariant-type pseudometrics.

2.4.3.3 Vanishing and non-degeneracy of fragmentation metrics. We fix here a TPC denoted by $\mathcal{C}$ together with the associated weights and pseudometrics, as above. We will denote by $\bar{d}^{\{0\}}$ the pseudometric associated to the family consisting of only the element 0. In view of point (iii) above, $\bar{d}^{\{0\}}$ is an upper bound for all the pseudometrics $\bar{d}^{\mathcal{F}}$.

It is obvious, as noticed in Remark 2.59, that in general $\bar{d}^{\mathcal{F}}$ is degenerate. For instance, if $\mathcal{F} = \mathrm{Obj}(\mathcal{C})$ then $\bar{d}^{\mathcal{F}} \equiv 0$. The rest of Remark 2.59 also continues to apply to $\bar{d}^{\mathcal{F}}$. We list below some other easily proven properties. We assume for all the objects X involved here that $\sigma(\mathbb{1}_X) = 0$, and we will use the calculations in Examples 2.70, 2.72, 2.74, and 2.75. Recall the notion of r-isomorphism from Definition 2.26; in particular, this is a morphism in $\mathcal{C}_0$. A 0-isomorphism is simply an isomorphism in the category $\mathcal{C}_0$ and is denoted by $\equiv$.

(i) If $X \equiv X'$, then $\bar{d}^{\mathcal{F}}(X, X') = 0 = \underline{\bar{d}}^{\mathcal{F}}(X, X')$ for any family $\mathcal{F}$.

(ii) We have

$$
\begin{aligned}
&2\,\bar{\delta}^{\{0\}}(X, X') \\
&\qquad \geq \inf\{r \in \mathbb{R} \mid \exists \text{ an } r\text{-isomorphism } \phi : \Sigma^k X' \to X \text{ for some } k \geq 0\} \\
&\qquad \geq \bar{\delta}^{\{0\}}(X, X').
\end{aligned}
$$

For the first inequality, consider a sequence of triangles (2.38) of total weight at most r. We intend to show that, for some $k \geq 0$, there exists a $2r$-isomorphism $\Sigma^k X' \to X$. We make use of Examples 2.77, 2.78, and 2.79. We assume, without loss of generality, that $T^{-1}X'$ appears in the jth triangle. The first triangles are of the form $0 \to Y_i \to Y_{i+1} \to 0$ for $i \leq j - 1$ with $Y_0 = 0$. By Example 2.78 we deduce that Y_{j-1} is r_0-acyclic and that r_0 is at most the sum of the weights of the first $j - 1$ triangles. The next triangle is of the form $T^{-1}X' \to Y_{j-1} \to Y_j \to X'$ and of weight r_j. Example 2.79 shows that there exists a $(2r_0 + r_j)$-isomorphism $X' \to \Sigma^{-k} Y_j$. The next triangles, of the form $0 \to Y_i \to Y_{i+1} \to 0$, have weights r_i, and there are r_i-isomorphisms $Y_i \to \Sigma^{-s_i} Y_{i+1}$ with $r_i \geq s_i$ (see Example 2.77). Putting things together, we have $2r \geq 2r' \geq 2r_0 + r_j + \cdots + r_n$,

and there is a $2r'$-isomorphism $X' \to \Sigma^{-s_1-s_2-s_3-\cdots-s_n} X$. This implies that $2\,\bar{\delta}^{\{0\}} \geq \inf$. For the second inequality, assume that $\phi : X' \to \Sigma^{-k} X$ is an r-isomorphism. We need to construct a cone decomposition of $\overline{w}$-weight at most r. We first assume $k \leq r$. The first triangle is $T^{-1}X' \to 0 \to X'$; it is exact in $\mathcal{C}_0$ and of weight 0. The second triangle is $0 \to X' \to X \to 0$. The associated strict exact triangle is $(0 \to X' \to \Sigma^{-k} X \to 0, r)$, and it uses ϕ in an obvious way to compare with the exact triangle $0 \to X' \to X' \to 0$ in $\mathcal{C}_0$. So we are left with the case $k > r$. In this case, the first triangle is $T^{-1}X' \to 0 \to \Sigma^{k-r} X' \to X'$. Its $\overline{w}$-weight is null. The next triangle is $0 \to \Sigma^{k-r} X' \to X \to 0$, the associated strict exact triangle being $(0 \to \Sigma^{k-r} X' \to \Sigma^{-r} X \to 0, r)$, where ϕ is now used to compare with the exact triangle $0 \to \Sigma^{k-r} X' \to \Sigma^{k-r} X' \to 0$.

(iii) We have

$$\underline{\bar{\delta}}^{\{0\}}(X, X') = \inf\{r \in \mathbb{R} \mid \exists \text{ an } r\text{-isomorphism } \phi : \Sigma^k X' \to X \text{ with } r \geq k \geq 0\}.$$

This happens because the first triangle in the sequence (2.38) is $0 \to X' \to Y_1 \to 0$ and the next triangles are of the form $0 \to Y_i \to Y_{i+1} \to 0$. Each of them has a weight estimated by the numbers r_i for which there is an r_i-isomorphism $Y_i \to \Sigma^{-s_i} Y_{i+1}$ with $0 \leq s_i \leq r_i$, which shows the claim.

(iv) We have $\bar{d}^{\{0\}}(X, \Sigma^r X) = r$ for any $r \in \mathbb{R}$ (this follows from Examples 2.70 and 2.75).

(v) For the shift-invariant metric $\widehat{d^{\{0\}}}$ induced by $\bar{d}^{\{0\}}$ through the formula (2.41) we have $\widehat{d^{\{0\}}}(X, \Sigma^r X) = 0$ for all X and $r \in \mathbb{R}$.

Thus $\bar{d}^{\{0\}}$ is finite for pairs of objects that are isomorphic in $\mathcal{C}_\infty$ and $\bar{d}^{\{0\}}(X, X')$ is the optimal upper bound r such that there are s-isomorphisms in $\mathcal{C}$ with $s \leq r$, from some positive shift of X to X' and, similarly, from some positive shift of X' to X. To some extent, $\bar{d}^{\{0\}}$ can be viewed as an abstract analog of the interleaving distance in the theory of persistence modules (cf. [47, Section 1.3]). We explore the relation with interleaving in more detail in Section 2.4.3.4 (see also Proposition 2.105).

Remark 2.83. There is another pseudometric $\bar{d}^{\emptyset}$, which means that $\mathcal{F} = \emptyset$. In this case, by definition, this $\bar{d}^{\emptyset}$ is again an algebraic analog of the interleaving distance. We will not use this pseudometric later in the memoir, so we do not further discuss its properties here.

For $\underline{\bar{d}}^{\{0\}}$ there is the additional constraint that the respective shifts should also be bounded by r. As a consequence:

(vi) If $\bar{d}^{\{0\}}(X, X') = 0$, then X and X' are 0-isomorphic up to shift. Moreover, if X and X' are not 0-isomorphic, they are both periodic in the sense

that there exist k and k' (not both null) and 0-isomorphisms $\Sigma^k X \to X$, $\Sigma^{k'} X' \to X'$.

(vii) If $\underline{\bar{d}}^{\{0\}}(X, X') = 0$, then $X \equiv X'$.

In summary, this means that the best we can expect from the fragmentation pseudometrics is that they should be non-degenerate on the space of 0-isomorphism types. From now on, we will say that a fragmentation pseudometric is non-degenerate if this is the case. Assuming no periodic objects exist, the metric $\bar{d}^{\{0\}}$ is non-degenerate in this sense. However, the distance it measures for two objects that are not isomorphic in $\mathcal{C}_\infty$ is infinite. On the other hand, a metric such as $\bar{d}^{\mathcal{F}}$ (as well as $\underline{\bar{d}}^{\mathcal{F}}$), where $\mathcal{F}$ is a family of triangular generators of $\mathcal{C}_\infty$, is finite but is in general degenerate.

The last point we want to raise in this section is that mixing fragmentation pseudometrics can sometimes produce non-degenerate ones. We will see an example of this sort in the symplectic Section 3.1, but we end here by describing a more general, abstract argument. Fix two families $\mathcal{F}_i$, $i = 1, 2$, of generators of $\mathcal{C}_\infty$. Consider the mixed pseudometric defined by

$$\bar{d}^{\mathcal{F}_1, \mathcal{F}_2} = \max\{\bar{d}^{\mathcal{F}_1}, \bar{d}^{\mathcal{F}_2}\}. \tag{2.42}$$

The idea is that if these two families are "separated" in a strong sense, then the mixed metric is non-degenerate. For instance, denote by $\mathcal{F}_i^\Delta$ the subcategory of $\mathcal{C}_0$ that is generated by $\mathcal{F}_i$. Now assume that $\mathrm{Obj}(\mathcal{F}_1^\Delta) \cap \mathrm{Obj}(\mathcal{F}_2^\Delta) = \{0\}$ (this is of course quite restrictive). We now claim that $\underline{\bar{d}}^{\mathcal{F}_1, \mathcal{F}_2}$ is non-degenerate and that $\bar{d}^{\mathcal{F}_1, \mathcal{F}_2}$ satisfies a weaker non-degeneracy condition, namely, that $\bar{d}^{\mathcal{F}_1, \mathcal{F}_2}(X, 0) = 0$ if and only if $X \equiv 0$. This latter fact follows immediately by noticing that $\bar{d}^{\mathcal{F}_i}(X, 0) = 0$ means that $X \in \mathcal{F}_i^\Delta$. We leave the former as an exercise.

2.4.3.4 Fragmentation pseudometrics, the interleaving pseudometric, and other algebraic measurements. The aim of this section is to describe relations between the fragmentation pseudometrics introduced above and the interleaving distance, which is well known in persistence theory as well as in Morse and Floer theory. A discussion of the bottleneck distance, which is closely related to interleaving, is included in Section 2.5.2.

Adapting the definition of the interleaving distance [47] to TPCs is immediate.

Definition 2.84. Let $\mathcal{C}$ be a triangulated persistence category with shift functor Σ. Given two objects $X, Y \in \mathrm{Obj}(\mathcal{C})$, the *interleaving distance* between X and Y is defined by

$$d_{\mathrm{int}}(X, Y) = \inf\big\{ r \geq 0 \mid \exists\, \phi \in \hom_{\mathcal{C}_0}(\Sigma^r X, Y),\ \psi \in \hom_{\mathcal{C}_0}(\Sigma^r Y, X) \\ \text{such that } \psi \circ \Sigma^r \phi = \eta^X_{2r},\ \phi \circ \Sigma^r \psi = \eta^Y_{2r} \big\}.$$

It is a simple exercise to check that this is indeed a pseudometric.

We will also make use of the shift-invariant version:

$$\widehat{d_{\text{int}}}(X,Y) = \inf_{r\in\mathbb{R}} d_{\text{int}}(\Sigma^r X, Y),$$

which we will refer to as the *shift-invariant interleaving pseudometric*.

Lemma 2.85. *Let $\mathcal{C}$ be a triangulated persistence category and let X, Y be two objects in $\mathcal{C}$. If $\phi : X \to Y$ is an r-isomorphism, then $d_{\text{int}}(X,Y) \leq r$. Conversely, if $d_{\text{int}}(X,Y) = s$, then for any $r > s$ there exist $4r$-isomorphisms $\Sigma^r X \to Y$ and $\Sigma^r Y \to X$.*

Proof. We start with the first part of the lemma. By the results in Section 2.3.1, ϕ has a right r-inverse $\psi : \Sigma^r Y \to X$ such that $\phi \circ \psi = \eta_r^Y$. Let $\phi' = \phi \circ \eta_r^X$ be such that we have the diagram

$$\Sigma^{2r} Y \xrightarrow{\Sigma^r \psi} \Sigma^r X \xrightarrow{\phi'} Y$$

with $\phi' \circ \Sigma^r \psi = \eta_{2r}^Y$. We now consider the composition $\psi \circ \Sigma^r \phi'$:

$$\Sigma^{2r} X \xrightarrow{\eta_r^{\Sigma^r X}} \Sigma^r X \xrightarrow{\Sigma^r \phi} \Sigma^r Y \xrightarrow{\psi} X.$$

Proposition 2.28 claims that ψ is r-equivalent to a left inverse of ϕ. Thus we have $\psi \circ \Sigma^r \phi \simeq_r \eta_r^X$, and hence $\psi \circ \Sigma^r \phi' = \eta_{2r}^X$, which shows the claim.

We pass to the second part of the lemma and now assume that $d_{\text{int}}(X,Y) = s < r$. We fix morphisms $f : \Sigma^r X \to Y$ and $g : \Sigma^r Y \to X$ such that $g \circ \Sigma^r f = \eta_{2r}^X$ and $f \circ \Sigma^r g = \eta_{2r}^Y$. This means that both f and g have right and left $2r$-inverses. Therefore, by Lemma 2.32 (iv), they are both $4r$-isomorphisms. ∎

Recall now the largest of our metrics from Section 2.4.3.1, $\bar{d}^{\{0\}}$, and its shift-invariant version $\widehat{d^{\{0\}}}$; see also Sections 2.4.3.2 and 2.4.3.3.

Corollary 2.86. *In the setting above we have*

$$\frac{1}{2}\,\widehat{d_{\text{int}}}(X,Y) \leq \widehat{d^{\{0\}}}(X,Y) \leq 4\,\widehat{d_{\text{int}}}(X,Y).$$

In particular, all shift-invariant pseudometrics of the type $\widehat{\bar{d}^{\mathcal{F}}}$ have as upper bound

$$4 \cdot \textit{(the shift-invariant interleaving pseudometric)}.$$

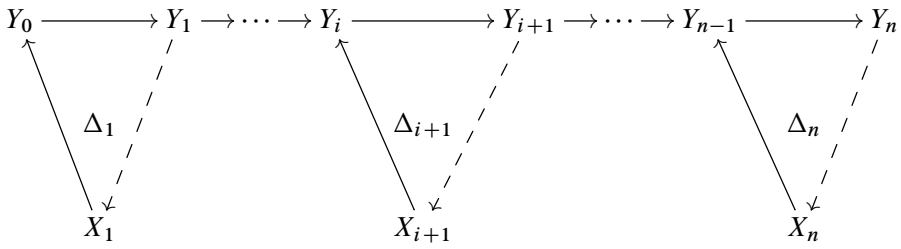
Proof. We start with the first inequality from the left. Assume that $\widehat{d^{\{0\}}}(X,Y) = s$ and fix $r > s$, $r < s + \epsilon$. This means that there exists $m \in \mathbb{R}$ such that $\bar{d}^{\{0\}}(\Sigma^m X, Y) < r$. In particular, $\bar{\delta}^{\{0\}}(\Sigma^m X, Y) < r$. Thus, by point (ii) in Section 2.4.3.3, we deduce that there exists some $k \geq 0$ and a $2r$-isomorphism $\Sigma^{m+k} X \to Y$. By the first point of Lemma 2.85 we deduce that $d_{\text{int}}(\Sigma^{m+k} X, Y) \leq 2r$, and thus $\widehat{d_{\text{int}}}(X,Y) \leq 2r$, which implies the desired inequality.

We now turn to the second inequality. Let $s = \widehat{d_{\mathrm{int}}}(X, Y)$, and let r be such that $s < r < s + \epsilon$. There exists $m \in \mathbb{R}$ such that $d_{\mathrm{int}}(\Sigma^m X, Y) < r$. From the second point of Lemma 2.85, we deduce that there are $4r$-isomorphisms $\Sigma^{r+m} X \to Y$ and $\Sigma^r Y \to \Sigma^m X$. This means, by point (ii) in Section 2.4.3.3, that $\bar{\delta}^{\{0\}}(\Sigma^m X, Y) \leq 4r$ and $\bar{\delta}^{\{0\}}(Y, \Sigma^m X) \leq 4r$. Thus $\bar{d}^{\{0\}}(\Sigma^m X, Y) \leq 4r$, and we conclude that $\widehat{d^{\{0\}}}(X, Y) \leq 4\widehat{d_{\mathrm{int}}}(X, Y)$. ∎

2.4.3.5 Other algebraic measurements. Other algebraic pseudometrics based on measuring the weight of cone decompositions – and not necessarily individual triangles – have appeared in [10]. The basic measurement introduced there can be viewed as a sort of extension of the interleaving distance and is easily formulated in the TPC setting (and in fact in any persistence category). To fix ideas, let $\mathcal{C}$ be a TPC and let $f : X \to Y$ be a morphism in $\mathcal{C}$. We define

$$\rho(f) = \inf_{g,s} \big\{ \max(\lceil f \rceil + \lceil g \rceil, s, 0) \mid g : Y \to X,\ g \circ f \simeq_s \mathbb{1}_X \big\}.$$

The way this is used in [10] is the following. Consider a triple (η, Y_n, ϕ) consisting of an iterated cone decomposition η in $\mathcal{C}_0$ with final term Y_n, as below:



and with $\phi : Y \to Y_n$ a morphism in $\mathcal{C}$ that induces an isomorphism in $\mathcal{C}_\infty$. The weight W of such a triple is defined by $W(\eta, Y_n, \phi) = \rho(\phi)$. This can be used to compare objects X, Y in $\mathcal{C}$ relative to a family of objects $\mathcal{F}$, which we assume to be closed under the action of Σ, by defining, for two objects X, Y,

$$\begin{aligned} &Z^{\mathcal{F}}(Y, X) \\ &\quad = \inf \big\{ W(\eta, Y_n, \phi) \mid \exists j, \forall i \neq j,\ X_i \in \mathcal{F},\ T^{-1} X = X_j,\ \phi : \Sigma^k Y \to Y_n,\ k \in \mathbb{R} \big\}. \end{aligned}$$

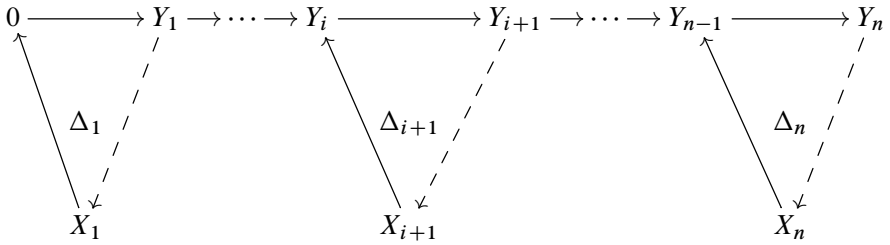
Such a $Z^{\mathcal{F}}$ can be obviously symmetrized and it is shift invariant (because $\mathcal{F}$ is closed under the action of Σ and we included the parameter k in the infimum). However, the fact that the triangle inequality is satisfied is non-trivial, and it is not clear whether this is true for general TPCs. As we will see later in Remark 2.103, the triangle inequality holds in important examples such as the homotopy category of a filtered pre-triangulated dg-category, and similarly for filtered modules over a filtered A_∞-category (this is the case treated in [10]). This subtlety is related to the degree of

precision in constructing maps induced on cones. Nonetheless, there is a simple way to compare some of our pseudometrics and $Z^{\mathcal{F}}$.

Lemma 2.87. *In the setting above we have*

$$Z^{\mathcal{F}}(Y,X) \leq 4 \cdot \widehat{\bar{\delta}^{\mathcal{F}}}(Y,X).$$

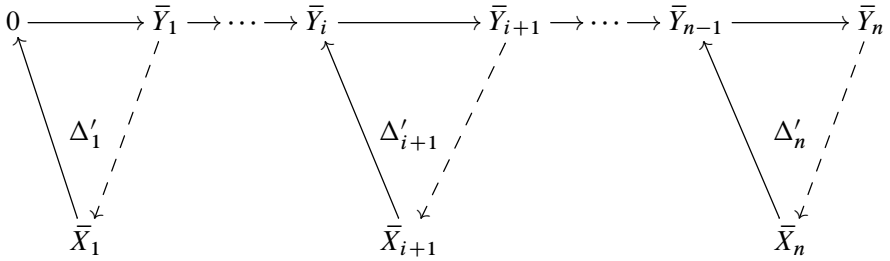
Proof. Assume that $\widehat{\bar{\delta}^{\mathcal{F}}}(Y,X) = r$. This means that there exists $k \in \mathbb{R}$ and a sequence of exact triangles η_∞ in $\mathcal{C}_\infty$



(2.43)

with $X_i \in \mathcal{F}$ for all indices except for X_j, in which case $X_j = \Sigma^l T^{-1} X$ for some $l \in \mathbb{R}$. Moreover, $Y_n = Y$ and $\bar{w}(\Delta_i) = r_i$ with $\sum_i r_i = r + \epsilon$ (for any small ϵ).

The aim is to construct another sequence $\bar{\eta}$ of exact triangles, this time exact in $\mathcal{C}_0$,



(2.44)

such that for each i there is a $k_i \in \mathbb{R}$ with $\overline{X}_i = \Sigma^{k_i} X_i$ and there is some $\phi : Y \to \overline{Y}_n$ that is a $2(r+\epsilon)$-isomorphism. Assuming this construction is achieved, we deduce from Lemma 2.85 that Y and $\overline{Y}_n$ are $2(r+\epsilon)$-interleaved and thus $\rho(\phi) \leq 4(r+\epsilon)$, which implies the claim.

The first step is to replace the sequence η_∞ with a sequence η of strict exact triangles in $\mathcal{C}$:

$$\begin{cases} \Delta_1 : X'_1 \to 0 \to Y'_1 \to \Sigma^{-r_1} T X'_1 \\ \Delta_2 : X'_2 \to Y'_1 \to Y'_2 \to \Sigma^{-r_2} T X'_2 \\ \Delta_3 : X'_3 \to Y''_2 \to Y'_3 \to \Sigma^{-r_3} T X'_3 \\ \qquad\qquad \vdots \\ \Delta_n : X'_n \to Y'_{n-1} \to Y'_n \to \Sigma^{-r_n} T X'_n, \end{cases}$$

where the $(-)'$ are appropriate shifts of the corresponding nodes of the sequence η_∞. Such triangles exist by the definition of the weight $\overline{w}$. We now proceed by induction: we assume that we have constructed the first k triangles as in $\bar{\eta}$ together with an $s_k = 2(r_1 + r_2 + \cdots + r_k)$-isomorphism $\phi_k : Y'_k \to \overline{Y}_k$. We now consider the strict exact triangle Δ_{k+1}:

$$\begin{array}{rccccc} \Delta_{k+1}: & X'_{k+1} & \xrightarrow{u_{k+1}} & Y'_k & \longrightarrow & Y''_{k+1} \longrightarrow TX'_{k+1} \\ & & & & & \downarrow f \\ & & & & & Y'_{k+1} \longrightarrow \Sigma^{-r_{k+1}} TX'_{k+1} \end{array}$$

as in Definition 2.42 with f an r_{k+1}-isomorphism and the top row an exact triangle in $\mathcal{C}_0$. We consider the two exact triangles in $\mathcal{C}_0$,

$$\begin{array}{ccccccc} X'_{k+1} & \xrightarrow{u_{k+1}} & Y'_k & \longrightarrow & Y''_{k+1} & \longrightarrow & TX'_{k+1} \\ \downarrow & & \downarrow{\scriptstyle \phi_k} & & \downarrow{\scriptstyle h} & & \downarrow \\ X'_{k+1} & \xrightarrow{u'} & \overline{Y}_k & \longrightarrow & Y'''_{k+1} & \longrightarrow & TX'_{k+1} \end{array} \tag{2.45}$$

where $u' = \phi_k \circ u_{k+1}$ and h is induced from the first square on the left. In particular, h is an s_k-isomorphism. So now we consider

$$Y'_{k+1} \xrightarrow{g} \Sigma^{-r_{k+1}} Y''_{k+1} \xrightarrow{h} \Sigma^{-r_{k+1}} Y'''_{k+1},$$

where g is a left r_{k+1}-inverse of f, and we notice that $h \circ g$ is an s_{k+1}-isomorphism. We will take the map ϕ_{k+1} to be the composition $\phi_{k+1} = h \circ g$ and we put $\overline{Y}_{k+1} = \Sigma^{-r_{k+1}} Y'''_{k+1}$. We take the triangle

$$\Sigma^{-r_{k+1}} X'_{k+1} \xrightarrow{u'} \Sigma^{-r_{k+1}} \overline{Y}_k \longrightarrow \Sigma^{-r_{k+1}} Y'''_{k+1} \longrightarrow \Sigma^{-r_{k+1}} TX'_{k+1},$$

which is the bottom row in (2.45) shifted by $\Sigma^{-r_{k+1}}$ as the $(k+1)$st exact triangle in the sequence $\bar{\eta}$. Finally, we adjust the first k triangles already constructed by shifting them all down by $\Sigma^{-r_{k+1}}$. This produces a sequence of $k+1$ triangles, each exact in $\mathcal{C}_0$, with the desired properties, together with the map ϕ_{k+1}; this completes the induction step. ∎

Possibly more useful than the actual statement of Lemma 2.87 is the method of proof: we produced a sequence $\bar{\eta}$ of exact triangles in $\mathcal{C}_0$, as in (2.44), and a $2r$-isomorphism $\phi : Y_n \to \overline{Y}_n$ out of the sequence of triangles in $\mathcal{C}_\infty$ in (2.43) whose sum of weights is r.

Using a right inverse $\psi : \Sigma^{2r} \overline{Y}_n \to Y_n$ of ϕ we can transform the last $\mathcal{C}_0$ exact triangle into a strict exact triangle of weight $4r$. The interest of this construction –

and this will be used in the applications in Section 3.1 – is that we obtain in this way a method to bound $\widehat{\bar{\delta}^{\mathcal{F}}}(Y,X)$ both from below and from above by a simpler quantity $Q^{\mathcal{F}}(Y,X)$, defined as the infimum of the sum of the weights of the triangles in decompositions as in (2.43), but with the first $n-1$ triangles of weight 0. Thus the weight of such a decomposition equals the weight of Δ_n. To summarize what was discussed above we have:

Corollary 2.88. *For any objects X, Y, we have*

$$\widehat{\bar{\delta}^{\mathcal{F}}}(Y,X) \leq Q^{\mathcal{F}}(Y,X) \leq 4 \cdot \widehat{\bar{\delta}^{\mathcal{F}}}(Y,X).$$

Remark 2.89. If in the inequality above one could avoid the factor 4, then we would have a simpler description of the fragmentation pseudometrics discussed here by replacing sequences of strict exact triangles in $\mathcal{C}$ with corresponding sequences in $\mathcal{C}_0$, followed by an s-isomorphism, where s is the sum of the weights of the initial triangles. However, this coefficient has to do with the fact that left (or right) inverses of k-isomorphisms are, in general, only $2k$-isomorphisms (see also Remark 2.34), and a factor of at least 2 is basically unavoidable.

2.5 Examples

2.5.1 Filtered dg-categories

The key property of dg-categories, introduced in [12] (see also [27]), is that they admit natural, pre-triangulated closures. The 0-cohomological category of this closure is triangulated. We will see here that there is a natural notion of *filtered* dg-categories. Such a category also admits a pre-triangulated closure, defined using filtered twisted complexes, following closely [12]. Its 0-cohomological category is a triangulated persistence category.

2.5.1.1 Basic definitions. Following a standard convention, we work in a co-homological setting and we keep all the sign conventions as in [12]. For our purposes it is convenient to view a filtered cochain complex over the field $\mathbf{k}$ as a triple (X, ∂, ℓ) consisting of a cochain complex (X, ∂) and a filtration function $\ell : X \to \mathbb{R} \cup \{-\infty\}$ such that for any $a, b \in X$ and $\lambda \in \mathbf{k} \setminus \{0\}$, $\ell(\lambda a + b) \leq \max\{\ell(a), \ell(b)\}$, $\ell(a) = -\infty$ if and only if $a = 0$, and $\ell(\partial a) \leq \ell(a)$. We denote by $X^{\leq r} = \{x \in X \mid \ell(x) \leq r\} \subset X$ the filtration induced on X by the filtration function ℓ. Clearly, $X^{\leq r}$ is again a filtered cochain complex. The family $\{X^{\leq r}\}_{r \in \mathbb{R}}$ determines the function ℓ. The cohomology of a filtered cochain complex is a persistence module: $V^r(X) = H(X^{\leq r}; \mathbf{k})$, whose structural maps $i_{r,s}$ are induced by the inclusions $\iota_{r,s} : X^{\leq r} \hookrightarrow X^{\leq s}$, $r \leq s$. We have omitted the grading here, as is customary. If it needs to be indicated, we write, for instance, $[V^r(X)]^i = H^i(X^{\leq r}; \mathbf{k})$. We denote this (graded) persistence

module by $\mathbb{V}(X)$:

$$\mathbb{V}(X) := (\{V^r(X)\}_{r\in\mathbb{R}}, \{i_{r,s} : V^r(X) \to V^s(X)\}_{r,s\in\mathbb{R}, r\le s}). \tag{2.46}$$

Given two filtered cochain complexes $X = (X, \partial^X, \ell_X)$ and $Y = (Y, \partial^Y, \ell_Y)$, their tensor product is a filtered cochain complex $(X \otimes Y, \partial^\otimes, \ell_\otimes)$ given by $(X \otimes Y)_k = \bigoplus_{i+j=k}(X_i \otimes Y_j)$ and

$$\partial^\otimes(x \otimes y) = \partial^X(x) \otimes y + (-1)^{|x|} x \otimes \partial^Y(y) \;\; \ell_\otimes(a \otimes b) = \ell_X(a) + \ell_Y(b).$$

If (X, ℓ_X) and (Y, ℓ_Y) are filtered vector spaces, we call a linear map $\phi : X \to Y$ *r-filtered* if $\ell_Y(\phi(x)) \le \ell_X(x) + r$ for all $x \in X$. A 0-filtered map is sometimes called (for brevity) filtered. For more background on this formalism, see [59].

The next definition is an obvious analog of the notion of dg-category in [12, Section 1].

Definition 2.90. A *filtered dg-category* is a preadditive category $\mathcal{A}$ where

(i) for any $A, B \in \mathrm{Obj}(\mathcal{A})$, the hom-set $\hom_{\mathcal{A}}(A, B)$ is a filtered cochain complex with filtrations denoted by $\hom^{\le r}_{\mathcal{A}}(A, B)$ such that, for each identity element, we have $\ell(\mathbb{1}_A) = 0$ and $\mathbb{1}_A$ is closed;

(ii) the composition is a filtered chain map:

$$\hom_{\mathcal{A}}(B, C) \otimes \hom_{\mathcal{A}}(A, B) \xrightarrow{\circ} \hom_{\mathcal{A}}(A, C);$$

(iii) for any inclusions $\iota^{AB}_{r,r'}$ and $\iota^{BC}_{s,s'}$, the composition morphism satisfies the compatibility condition $\iota^{BC}_{s,s'}(g) \circ \iota^{AB}_{r,r'}(f) = \iota^{AC}_{r+s,r'+s'}(g \circ f)$ for any $f \in \hom^{\le r}_{\mathcal{A}}(A, B)$ and $g \in \hom^{\le s}_{\mathcal{A}}(B, C)$.

Remark 2.91. A filtered dg-category is trivially a persistence category by forgetting the boundary maps on each $\hom_{\mathcal{A}}(A, B)$. Explicitly, for any $A, B \in \mathrm{Obj}(\mathcal{A})$, define $E_{AB} : (\mathbb{R}, \le) \to \mathrm{Vect}_{\mathbf{k}}$ by $E_{AB}(r) = \hom^{\le r}_{\mathcal{A}}(A, B)$ and $E_{AB}(i_{r,s}) = \iota_{r,s} : \hom^{\le r}_{\mathcal{A}}(A, B) \to \hom^{\le s}_{\mathcal{A}}(A, B)$.

The (co)homology category of a filtered dg-category $\mathcal{A}$, denoted by $\mathrm{H}(\mathcal{A})$, is a category with

$$\mathrm{Obj}(\mathrm{H}(\mathcal{A})) = \mathrm{Obj}(\mathcal{A})$$

and, for any $A, B \in \mathrm{Obj}(\mathrm{H}(\mathcal{A}))$,

$$\hom_{\mathrm{H}(\mathcal{A})}(A, B) := \mathbb{V}(\hom_{\mathcal{A}}(A, B)) = \left(\left\{H^*\left(\hom^{\le r}_{\mathcal{A}}(A, B)\right)\right\}_{r\in\mathbb{R}}, \{i_{r,s}\}_{r\le s}\right)$$

is the persistence module as described in (2.46). It is immediate to see that, for any filtered dg-category $\mathcal{A}$, its (co)homology category $\mathrm{H}(\mathcal{A})$ is a (graded) persistence category.

2.5.1.2 Twisted complexes. It is easy to construct a formal shift completion of a dg-category.

Definition 2.92. Let $\mathcal{A}$ be a filtered dg-category. The *shift completion* $\Sigma\mathcal{A}$ of $\mathcal{A}$ is a filtered dg-category such that

(i) the objects of $\Sigma\mathcal{A}$ are

$$\mathrm{Obj}(\Sigma\mathcal{A}) = \left\{\Sigma^r A[d] \mid A \in \mathrm{Obj}(\mathcal{A}),\, r \in \mathbb{R}, \text{ and } d \in \mathbb{Z}\right\},$$

where

$$\Sigma^0 A = A, \quad \Sigma^s(\Sigma^r A) = \Sigma^{r+s} A, \quad A[0] = A,$$
$$(A[d_1])[d_2] = A[d_1 + d_2], \quad (\Sigma^r A)[d] = \Sigma^r(A[d])$$

for any $r, s \in \mathbb{R}$ and $d_1, d_2, d \in \mathbb{Z}$;

(ii) for any $\Sigma^r A[d_A], \Sigma^s B[d_B] \in \mathrm{Obj}(\mathcal{A})$, the hom-set $\hom(\Sigma^r A[d_A], \Sigma^s B[d_B])$ is a filtered cochain complex with the same underlying cochain complex of $\hom(A, B)$ but with degree shifted by $d_B - d_A$ and filtration function $\ell_{\Sigma^r A[d_A]\,\Sigma^s B[d_B]} = \ell_{AB} + s - r$.

Remark 2.93. It is immediate to check that $\Sigma\mathcal{A}$ as given in Definition 2.92 is still a filtered dg-category.

The category $\Sigma\mathcal{A}$ carries a natural functor $\Sigma : (\mathbb{R}, +) \to \mathcal{P}\mathrm{End}(\Sigma\mathcal{A})$ defined on objects by $\Sigma^r(A) = \Sigma^r A$ and with an obvious definition on morphisms such that Σ^r is filtration preserving. For any $r, s \in \mathbb{R}$, the natural transformations $\eta_{r,s} : \Sigma^r \to \Sigma^s$ are such that $(\eta_{r,s})_A : \Sigma^r A \to \Sigma^s A$ is induced by the identity map $\mathbb{1}_A$ for each $A \in \mathrm{Obj}(\mathcal{A})$. In this context, we have a natural definition of (one-sided) twisted complexes obtained by adjusting [12, Section 4, Definition 1] to the filtered case.

Definition 2.94. Let $\mathcal{A}$ be a filtered dg-category. A *filtered (one-sided) twisted complex* of $\Sigma\mathcal{A}$ is a pair $A = \left(\bigoplus_{i=1}^n \Sigma^{r_i} A_i[d_i], q = (q_{ij})_{1\le i,j\le n}\right)$ such that the following conditions hold:

(i) $\Sigma^{r_i} A_i[d_i] \in \mathrm{Obj}(\Sigma\mathcal{A})$, where $r_i \in \mathbb{R}$ and $d_i \in \mathbb{Z}$.

(ii) $q_{ij} \in \hom_{\Sigma\mathcal{A}}(\Sigma^{r_j} A_j[d_j], \Sigma^{r_i} A_i[d_i])$ is of degree 1, and $q_{ij} = 0$ for $i \ge j$.

(iii) $d_{\hom} q_{ij} + \sum_{k=1}^n q_{ik} \circ q_{kj} = 0$.

(iv) For any q_{ij}, $\ell_{\Sigma^{r_j} A_j[d_j]\Sigma^{r_i} A_i[d_i]}(q_{ij}) \le 0$.

Remark 2.95. We will mostly work with *filtered* one-sided twisted complexes as defined above, but, more generally, the pair $A = \left(\bigoplus_{i=1}^n \Sigma^{r_i} A_i[d_i], q = (q_{ij})_{1\le i,j\le n}\right)$, subject only to (i), (ii), and (iii), is called a one-sided twisted complex.

It is easy to see that there are at least as many filtered one-sided twisted complexes as one-sided twisted complexes; this follows from the statement below, whose proof we leave to the reader.

Lemma 2.96. *Given a twisted complex* $\left(\bigoplus_{i=1}^{n} A_i[d_i], q = (q_{ij})\right)$*, there exist* $(r_i)_{1\leq i\leq n}$ *such that condition (iv) in Definition* 2.94 *is satisfied for the filtration-shifted twisted complex* $\left(\bigoplus_{i=1}^{n} \Sigma^{r_i} A_i[d_i], q = (q_{ij})\right)$.

2.5.1.3 Pre-triangulated completion. We will see next that the filtered twisted complexes over $\mathcal{A}$ form a category that provides a (pre-)triangulated closure of $\mathcal{A}$. The 0-cohomology category of this completion is a triangulated persistence category.

Definition 2.97. Given a filtered dg-category $\mathcal{A}$, define its *filtered pre-triangulated completion*, denoted by $\mathrm{Tw}(\mathcal{A})$, to be the category with the following properties:

(i) Its objects are

$$\mathrm{Obj}(\mathrm{Tw}(\mathcal{A})) := \{\text{filtered one-sided twisted complexes of } \Sigma\mathcal{A}\}.$$

(ii) For $A = \left(\bigoplus \Sigma^{r_j} A_j[d_j], q\right)$ and $A' = \left(\bigoplus \Sigma^{r'_i} A'_i[d'_i], q'\right)$ in $\mathrm{Obj}(\mathrm{Tw}(\mathcal{A}))$, a morphism $f \in \hom_{\mathrm{Tw}(\mathcal{A})}(A, A')$ is a matrix of morphisms in $\mathcal{A}$, denoted by $f = (f_{ij}) : A \to A'$, where

$$f_{ij} \in \hom_{\Sigma\mathcal{A}}\left(\Sigma^{r_j} A_j[d_j], \Sigma^{r'_i} A'_i[d'_i]\right).$$

(iii) The hom-differential is defined as follows. For any $f \in \hom_{\mathrm{Tw}(\mathcal{A})}(A, A')$ as in (ii) above, define

$$d_{\mathrm{Tw}(\mathcal{A})}(f) := (d_{\hom} f_{ij}) + q' f - (-1)^l f q,$$

where $\deg(f_{ij}) = l$, and the right-hand side is written in matrix form. The composition $f' \circ f$ is given by matrix multiplication.

Lemma 2.98. *Given a filtered dg-category* $\mathcal{A}$*, its filtered pre-triangulated completion* $\mathrm{Tw}(\mathcal{A})$ *is a filtered dg-category.*

Proof. The key step is to note that there exists a filtration function on $\hom_{\mathrm{Tw}(\mathcal{A})}(A, A')$ for any $A, A' \in \mathrm{Obj}(\mathrm{Tw}(\mathcal{A}))$. For any $f = (f_{ij}) \in \hom_{\mathrm{Tw}(\mathcal{A})}(A, A')$, set

$$\ell_{AA'}(f) = \max_{i,j}\left\{\ell_{\Sigma^{r_j} A_j[d_j]\,\Sigma^{r'_i} A'_i[d'_i]}(f_{ij})\right\}.$$

It is easily checked that $\ell_{AA'}$ is a filtration function and that it satisfies the other required properties. ∎

The first step towards triangulation is to define an appropriate cone of a morphism.

Definition 2.99. Let $\mathcal{A}$ be a filtered dg-category and $\mathrm{Tw}(\mathcal{A})$ be its pre-triangulated completion. Let

$$A = \left(\bigoplus \Sigma^{r_j} A_j[d_j], q = (q_{ij})_{1\le i,j\le n}\right),$$
$$A' = \left(\bigoplus \Sigma^{r'_i} A'_i[d'_i], q = (q'_{ij})_{1\le i,j\le m}\right)$$

be two objects of $\mathrm{Tw}(\mathcal{A})$, and let $f : A \to A'$ be a closed, degree-preserving morphism. Define the λ-*filtered mapping cone* of f, where $\lambda \ge \ell_{AA'}(f)$, by

$$\mathrm{Cone}^{\lambda}(f) := \left(\bigoplus_i \Sigma^{r'_i} A'_i[d'_i] \oplus \bigoplus_j \Sigma^{r_j+\lambda} A_j[d_j+1], q_{\mathrm{co}}\right),$$

where $q_{\mathrm{co}} = \begin{pmatrix} q' & f \\ 0 & -q \end{pmatrix}$, with q', q, f all block matrices.

Remark 2.100. (a) The condition $\lambda \ge \ell_{AA'}(f)$ guarantees that $\mathrm{Cone}^{\lambda}(f)$ is indeed a filtered one-sided twisted complex over $\Sigma\mathcal{A}$. Therefore, $\mathrm{Tw}(\mathcal{A})$ is closed under taking degree shifts, filtration shifts, and filtered mapping cones of (degree-preserving) closed morphisms.

(b) Notice that a λ-filtered cone can also be written as a 0-filtered cone but for a different map.

(c) Given a filtered dg-category $\mathcal{A}$, it is easy to see that every object in $\mathrm{Tw}(\mathcal{A})$ can be obtained from objects in $\Sigma\mathcal{A}$ by taking iterated 0-filtered mapping cones.

The 0-cohomological category associated to a dg-category is a triangulated category. The next result is the analog in the filtered case.

Proposition 2.101. *If $\mathcal{A}$ is a filtered dg-category and $\mathrm{Tw}(\mathcal{A})$ is its filtered pre-triangulated completion, then the degree-0 cohomology category $H^0(\mathrm{Tw}(\mathcal{A}))$ is a triangulated persistence category.*

In view of this result, it is natural to call a filtered dg-category $\mathcal{A}$ *pre-triangulated* if the inclusion $\mathcal{A} \hookrightarrow \mathrm{Tw}(\mathcal{A})$ is an equivalence of filtered dg-categories.

Corollary 2.102. *Let $\mathcal{A}$ be a filtered pre-triangulated dg-category. Then its degree-0 cohomology category $\mathrm{H}^0(\mathcal{A})$ is a triangulated persistence category.*

Proof of Proposition 2.101. It is trivial to notice that the category $H^0(\mathrm{Tw}(\mathcal{A}))$ is a persistence category. It is endowed with an obvious shift functor as defined in Section 2.5.1.2. The first thing to check is that the 0-level category $[H^0(\mathrm{Tw}(\mathcal{A}))]_0$, with the same objects as $H^0(\mathrm{Tw}(\mathcal{A}))$ and only the shift 0-morphisms, is triangulated; see Definition 2.21. The family of triangles that will provide the exact ones are the triangles of the form

$$A \xrightarrow{f} B \xrightarrow{i} \mathrm{Cone}^0(f) \xrightarrow{\pi} A[1]$$

associated to the 0-cones, as given in Definition 2.99. From this point on, checking that $[H^0(\mathrm{Tw}(\mathcal{A}))]_0$ is triangulated comes down to the usual verifications showing that the H^0 of a dg-category is triangulated, with a bit of care to ensure that the relevant homotopies preserve filtration. We leave this verification to the reader. It is then automatic that Σ^r is triangulated when restricted to $[H^0(\mathrm{Tw}(\mathcal{A}))]_0$. The last step is to show that the morphism $\eta_r^A : \Sigma^r A \to A$ has an r-acyclic cone in $\mathrm{Tw}(\mathcal{A})$. In this context of filtered dg-categories, an object K is r-acyclic if the identity $\mathbb{1}_K \in \hom_{\mathrm{Tw}(\mathcal{A})}(K, K)$ is a boundary of some element $\eta \in \hom^{\leq r}_{\mathrm{Tw}(\mathcal{A})}(K, K)$.

The map $\eta_r^A \in \mathrm{Mor}^0_{\mathrm{Tw}(\mathcal{A})}(\Sigma^r A, A)$ is induced by the identity. By definition, we have $\mathrm{Cone}^0(\eta_r^A) = A \oplus \Sigma^r A[1]$ and

$$q_{\mathrm{co}} = \begin{pmatrix} q & \eta_r^A \\ 0 & -q' \end{pmatrix},$$

where q is the structural map of the twisted complex A and $q' = \Sigma^r q$. Consider a homotopy

$$K = \begin{pmatrix} 0 & 0 \\ (\eta_{0,r})_A & 0 \end{pmatrix} : \mathrm{Cone}^0(\eta_r^A) \to \mathrm{Cone}^0(\eta_r^A)[1].$$

Note that $\ell(K) = r$. We have

$$\begin{aligned} dK &= \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} q & \eta_r^A \\ 0 & -q' \end{pmatrix} \begin{pmatrix} 0 & 0 \\ (\eta_{0,r})_A & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ (\eta_{0,r})_A & 0 \end{pmatrix} \begin{pmatrix} q & \eta_r^A \\ 0 & -q' \end{pmatrix} \\ &= \begin{pmatrix} \mathbb{1}_A & 0 \\ -q' \circ (\eta_{0,r})_A + (\eta_{0,r})_A \circ q & \mathbb{1}_{\Sigma^r A[1]} \end{pmatrix} = \mathbb{1}_{\mathrm{Cone}^0(\eta_r^A)}, \end{aligned}$$

because $\Sigma^r q \circ (\eta_{0,r})_A = (\eta_{0,r})_A \circ q$ and this concludes the proof. ∎

Remark 2.103. (a) In the filtered dg-category $\mathrm{Tw}(\mathcal{A})$, we can replicate all the constructions in Section 2.3 at the chain level, similarly to the definition of r-acyclic objects mentioned in the proof above. For instance, r-isomorphisms are replaced by r-quasi-isomorphisms (meaning filtration-preserving morphisms that induce an r-isomorphism in homology), and all the functorial-type constructions of that section can be pursued at the chain level, by replacing commutativity at the chain level by commutativity up to homotopy.

(b) One advantage of working at the chain level instead of in the general setting of triangulated persistence categories is that the maps c induced on cones through diagrams of the form

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \longrightarrow & \mathrm{Cone}(f) \\ \downarrow{\scriptstyle a} & & \downarrow{\scriptstyle b} & & \downarrow{\scriptstyle c} \\ A' & \xrightarrow{f'} & B' & \longrightarrow & \mathrm{Cone}(f') \end{array}$$

are defined explicitly in terms of the homotopy making the square on the left commutative. An example relevant for this memoir is that, in the homological category of a filtered dg-category, the measurement $Z^{\mathcal{F}}(-,-)$ from Section 2.4.3.5 satisfies the triangle inequality. The proof follows closely the arguments in [10, Lemma 6.11], with all weakly filtered maps there being replaced with filtered ones here.

2.5.2 Filtered cochain complexes

In this section we discuss the main example of a filtered dg-category, the category of filtered co-chain complexes. As we shall see, this is pre-triangulated and thus, in view of Corollary 2.102, its homotopy category is a triangulated persistence category.

We will work over a field $\mathbf{k}$ and will denote the resulting category by $\mathcal{FK}_{\mathbf{k}}$. The objects of this category are filtered cochain complexes (X, ∂, ℓ), where (X, ∂) is a cochain complex and ℓ is a filtration function, as in Section 2.5.1.1. Given two filtered cochain complexes (X, ∂_X, ℓ_X) and (Y, ∂_Y, ℓ_Y), the morphisms $\hom_{\mathcal{FK}_{\mathbf{k}}}(X, Y)$ are linear graded maps $f : X \to Y$ such that the quantity

$$\ell(f) = \inf\{r \in \mathbb{R} \mid \ell_Y(f(x)) \leq \ell_X(x) + r \ \forall x \in X\} \tag{2.47}$$

is finite. The filtration function on $\hom_{\mathcal{FK}_k}(X, Y)$ is then defined through (2.47). The differential on $\hom_{\mathcal{FK}_k}(X, Y)$ is given, as usual, by $\partial(f) = \partial_Y \circ f - (-1)^{|f|} f \circ \partial_X$, and it obviously preserves filtrations. The composition of morphisms is also obviously compatible with the filtration and therefore $\mathcal{FK}_{\mathbf{k}}$ is a filtered dg-category.

There is a natural shift functor on $\mathcal{FK}_{\mathbf{k}}$, $\Sigma : (\mathbb{R}, +) \to \mathcal{P}\mathrm{End}(\mathcal{F}\mathcal{K}_k)$, defined by

$$\Sigma^r(X, \partial, \ell_X) = (X, \partial, \ell_X + r)$$

and

$$\Sigma^r(f) = f \quad \text{for any } f \in \mathrm{Mor}_{\mathcal{FK}_k}(X, Y).$$

Moreover, for $r, s \in \mathbb{R}$, there is a natural transformation from Σ^r to Σ^s induced by the identity.

Assume that $f : (X, \partial_X, \ell_X) \to (Y, \partial_Y, \ell_Y)$ is a cochain morphism such that $\ell(f) \leq 0$. In this case, the usual cone construction $\mathrm{Cone}(f) = (Y \oplus X[1], \partial_{\mathrm{co}})$, with

$$\partial_{\mathrm{co}} = \begin{pmatrix} \partial_Y & f \\ 0 & -\partial_X \end{pmatrix},$$

produces a filtered complex and fits into a triangle of maps with $\ell \leq 0$:

$$X \xrightarrow{f} Y \xrightarrow{i} \mathrm{Cone}(f) \xrightarrow{\pi} X[1].$$

The standard properties of this construction immediately imply that the dg-category $\mathcal{FK}_{\mathbf{k}}$ is pre-triangulated and thus the 0-cohomological category, $H^0\mathcal{FK}_{\mathbf{k}}$, is a triangulated persistence category.

It is useful to make explicit some of the properties of this category:

(i) The objects of $H^0\mathcal{FK}_{\mathbf{k}}$ are filtered cochain complexes (X, ∂_X, ℓ_X).

(ii) The morphisms in $\hom^r_{H^0\mathcal{FK}_{\mathbf{k}}}(X, Y)$ are cochain maps $f : (X, \partial_X, \ell_X) \to (Y, \partial_Y, \ell_Y)$ such that $\ell(f) \leq r$ up to chain homotopy $h : f \simeq f'$ with $\ell(h) \leq r$.

(iii) A filtered complex (K, ∂_K, ℓ_K) is r-acyclic if the identity $\mathbb{1}_K$ is chain homotopic to 0 through a chain homotopy $h : \mathbb{1}_K \simeq 0$ with $\ell(h) \leq r$.

(iv) The construction of weighted exact triangles, as well as their properties, can be pursued in this context by following closely the scheme in Section 2.3.3.

(v) The limit category $[H^0\mathcal{FK}_{\mathbf{k}}]_\infty$ has as morphisms chain homotopy classes of cochain maps (where both the cochain maps and the homotopies are assumed to be of bounded shifts). Its objects are still filtered cochain complexes. It is triangulated, with translation functor $TX = X[1]$, as expected.

Remark 2.104. The example of the dg-category $\mathcal{FK}_{\mathbf{k}}$ can be extended in a number of ways, and we mention a couple of them here.

(a) Assume that we fix a filtered dg-category $\mathcal{A}$. There is a natural notion of a filtered (left/right) module $\mathcal{M}$ over $\mathcal{A}$. Such modules, together with filtered maps relating them, form a new filtered dg-category, denoted by $\mathrm{Mod}_{\mathcal{A}}$. The 0-cohomology category associated to this filtered dg-category, $H^0\mathrm{Mod}_{\mathcal{A}}$, is pre-triangulated because the category $\mathrm{Mod}_{\mathcal{A}}$ is naturally endowed with a shift functor, just like $\mathcal{FK}_{\mathbf{k}}$, as well as an appropriate cone construction over filtered, closed, degree-preserving morphisms.

(b) Similarly to (a), we may take $\mathcal{A}$ to be a filtered A_∞-category and consider the category $\mathrm{Mod}_{\mathcal{A}}$ of filtered modules over $\mathcal{A}$. Again this is a filtered dg-category and it is pre-triangulated (the formalism required to establish this fact appears in [10], in a version dealing with weakly filtered structures).

As mentioned at the beginning of Section 1.1, there exists a quantitative comparison between two filtered cochain complexes X, Y, called the bottleneck distance and denoted by $d_{\mathrm{bot}}(X, Y)$. This is best expressed in the barcode language from [3] or [59].

For completeness, we specify the version of barcodes used here. A *barcode* $\mathcal{B} = \{(I_j, m_j)\}_{j\in\mathcal{J}}$ is a collection of pairs consisting of intervals $I_j \subset \mathbb{R}$ and positive integers $m_j \in \mathbb{Z}_{>0}$, indexed by a set $\mathcal{J}$, and satisfying the following *admissibility* conditions:

- $\mathcal{J}$ is assumed to be either finite or equal to $\mathbb{Z}_{\geq 0}$.
- Each interval I_j is of the form $I_j = [a_j, b_j)$, with $-\infty < a_j < b_j \leq \infty$.
- In the case $\mathcal{J} = \mathbb{Z}_{\geq 0}$, we assume that $a_j \to \infty$ as $j \to \infty$.

The intervals I_j are called bars and, for each j, m_j is called the multiplicity of the bar I_j. To such a barcode one can associate a persistence module $V(\mathcal{B})$ that satisfies the following conditions:

- (Lower semi-continuity) For any $s \in \mathbb{R}$ and any $t \geq s$ sufficiently close to s, the map $i_{s,t} : M^s \to M^t$ is an isomorphism.
- (Lower boundedness) For s sufficiently small, $M^s = 0$.
- (Tameness) For every $s \in \mathbb{R}$, $\dim_{\mathbf{k}}(M^s) < \infty$.

The module $V(\mathcal{B})$ is defined as the direct sum of the elementary persistence modules $V(I)$ for each bar I in the barcode $\mathcal{B}$. Here $V([a,b))^s = \mathbf{k}$ if $s \in [a,b)$ and $V([a,b))^s = 0$ if $s \notin [a,b)$. Conversely, the Normal Form Theorem in [47, Section 2.1], or the main result in [25], says that any persistence module M with the three properties above can be decomposed as a direct sum of persistence modules of the form $V([a,b))$ and $V([c,\infty))$ in a unique way, up to permutation. Thus we can associate to it a barcode $\mathcal{B}(M)$ consisting of the intervals $[a,b)$ and $[c,\infty)$ appearing in the decomposition.

The homology $H(X)$ of a filtered cochain complex X is a persistence module whose barcode can be read out of the normal form of X. More precisely, by [59, Proposition 7.4] (see also [3]) there is a filtered isomorphism (in the category $H^0\mathcal{FK}_{\mathbf{k}}$) as follows:

$$X \simeq \bigoplus_{[a,+\infty)\in\mathcal{B}(X)} E_1(a) \oplus \bigoplus_{[c,d)\in\mathcal{B}(X)} E_2(c,d), \tag{2.48}$$

where $E_1(a), E_2(c,d) \in \mathrm{Obj}(\mathcal{FK}_{\mathbf{k}})$ are filtered cochain complexes defined by

$$E_1(a) = ((\cdots \to 0 \to \mathbf{k}\langle x\rangle \to 0 \to \cdots), \ell(x) = a)$$

and

$$E_2(c,d) = ((\cdots \to 0 \to \mathbf{k}\langle y\rangle \xrightarrow{\partial} \mathbf{k}\langle x\rangle \to 0 \to \cdots),\ \ell(y) = c, \ell(x) = d),$$

where $c \geq d$ and $\partial(y) = \kappa x$ for some $0 \neq \kappa \in \mathbf{k}$. The notation $\mathcal{B}(X)$ in (2.48) stands for a collection of intervals of two types: finite or semi-infinite intervals in $\mathbb{R}$ of the form $[c,d)$, with $c < d$ and possibly with $d = +\infty$; and intervals of length 0, $[c,d]$, with $c = d$.

In what follows, sometimes for brevity, we write $E_*(I)$ for $E_1(a)$ when $I = [a,+\infty)$, and for $E_2(c,d)$ when $I = [c,d)$, with $I \in \mathcal{B}(X)$. Then $d_{\mathrm{bot}}(X,Y)$ is defined as the infimum τ satisfying the following conditions: there exist some subsets consisting of certain "short intervals" $\mathcal{B}(X)_{\mathrm{short}} \subset \mathcal{B}(X)$ and $\mathcal{B}(Y)_{\mathrm{short}} \subset \mathcal{B}(Y)$ such that

(i) each short interval $[c,d)$ satisfies $2(d-c) \leq \tau$;

(ii) there is a bijection $\sigma : \mathcal{b}(x) \setminus \mathcal{b}(x)_{\mathrm{short}} \to \mathcal{b}(y) \setminus \mathcal{b}(y)_{\mathrm{short}}$;

(iii) if $\sigma([a,b)) = [c,d)$, then $\max\{|a-c|, |b-d|\} \le \tau$;

(iv) if $\sigma([a,\infty)) = [c,\infty)$, then $|a-c| \le \tau$.

In what follows, we assume that the cardinalities of the barcodes, $\#|\mathcal{B}(X)|$ and $\#|\mathcal{B}(Y)|$, are both finite. The following result compares the fragmentation pseudometric $d^{\mathcal{F}}$ defined in Definition 2.58 with the bottleneck distance d_{bot} defined above.

Proposition 2.105. *Let $\mathcal{C} = H^0\mathcal{FK}_{\mathbf{k}}$ and let $\mathcal{F} \subset \mathrm{Obj}(\mathcal{C})$ be a subset containing 0. Then*

$$d^{\mathcal{F}}(X,Y) \le C_{X,Y}\, d_{\mathrm{bot}}(X,Y),$$

where $C_{X,Y} = 4\min\{\#|\mathcal{B}(X)|, \#|\mathcal{B}(Y)|\} + 1$.

Proof. It is immediate to see that we may assume that both $\mathcal{B}'(X)$ and $\mathcal{B}'(Y)$ do not contain any 0-length bars, and thus $\mathcal{B}'(X) = \mathcal{B}(X)$, and similarly for Y. It suffices to prove the conclusion when $\mathcal{B}(X)$ and $\mathcal{B}(Y)$ have the same cardinality of the infinite-length bars (otherwise by definition $d_{\mathrm{bot}}(X,Y) = +\infty$ and the conclusion holds trivially). Let $\tau := d_{\mathrm{bot}}(X,Y) + \epsilon$ for an arbitrarily small $\epsilon > 0$. Since $d^{\mathcal{F}}(\cdot,\cdot)$ is invariant under filtered isomorphisms (applied to either of its two inputs), by (2.48) and by reordering summands we obtain

$$\begin{aligned} d^{\mathcal{F}}(X,Y) \le\ & d^{\mathcal{F}}\left(\bigoplus_{I \in \mathcal{B}(X)\setminus\mathcal{B}(X)_{\mathrm{short}}} E_*(I), \bigoplus_{\sigma(I)\in\mathcal{B}(Y)\setminus\mathcal{B}(Y)_{\mathrm{short}}} E_*(\sigma(I)) \right) \\ & + d^{\mathcal{F}}\left(\bigoplus_{J\in\mathcal{B}(X)_{\mathrm{short}}} E_2(J), \bigoplus_{J'\in\mathcal{B}(Y)_{\mathrm{short}}} E_2(J') \right), \end{aligned}$$

where the inequality is given by the triangle inequality of $d^{\mathcal{F}}$ with respect to the direct sum; see Proposition 2.60. For $d^{\mathcal{F}}$ with short intervals, both $\bigoplus_{J\in\mathcal{B}(X)_{\mathrm{short}}} E_2(J)$ and $\bigoplus_{J'\in\mathcal{B}(Y)_{\mathrm{short}}} E_2(J')$ are acyclic objects in $\mathcal{C}$; therefore by (i) in the definition of d_{bot} above, triangles

$$0 \to 0 \to \bigoplus_{J\in\mathcal{B}(X)_{\mathrm{short}}} E_2(J) \to 0 \quad \text{and} \quad 0 \to 0 \to \bigoplus_{J'\in\mathcal{B}(Y)_{\mathrm{short}}} E_2(J') \to 0$$

are weight-$\frac{\tau}{2}$ exact triangles (here we identify $\Sigma^\lambda 0$ with 0 for any shift $\lambda \in \mathbb{R}$). Thus,

$$\begin{aligned} d^{\mathcal{F}}\left(\bigoplus_{J\in\mathcal{B}(X)_{\mathrm{short}}} E_2(J), \bigoplus_{J'\in\mathcal{B}(X)_{\mathrm{short}}} E_2(J') \right) \le\ & d^{\mathcal{F}}\left(\bigoplus_{J\in\mathcal{B}(X)_{\mathrm{short}}} E_2(J), 0 \right) \\ & + d^{\mathcal{F}}\left(0, \bigoplus_{J'\in\mathcal{B}(X)_{\mathrm{short}}} E_2(J') \right) \le \tau. \end{aligned}$$

On the other hand, by Proposition 2.60 again, for $d^{\mathcal{F}}$ with non-short intervals, we have

$$d^{\mathcal{F}}\left(\bigoplus_{I\in\mathcal{B}(X)\setminus\mathcal{B}(X)_{\text{short}}} E_*(I), \bigoplus_{\sigma(I)\in\mathcal{B}(Y)} E_*(\sigma(I))\right) \leq \sum_{I\in\mathcal{B}(X)\setminus\mathcal{B}(X)_{\text{short}}} d^{\mathcal{F}}\big(E_*(I), E_*(\sigma(I))\big).$$

Since $d_{\text{bot}}(X,Y) < +\infty$, the bijection σ will always map a finite interval to a finite interval and a semi-infinite interval to a semi-infinite interval, so it suffices to consider the following two cases.

Case I. Estimate $d^{\mathcal{F}}(E_1(a), E_1(c))$. We need to build the desired cone decomposition. Without loss of generality, assume $a \geq c$. Then the identity map $\langle x\rangle_{E_1(a)} \to \langle x\rangle_{E_1(c)}$ (with *negative* filtration shift) implies that the triangle $E_1(a) \to E_1(c) \to K \to E_1(a)[1]$ is a weight-0 exact triangle (in fact in $\mathcal{C}_0$), where K is the filtered mapping cone. Then in the following cone decomposition (with linearization $(0, E_1(c))$):

$$\begin{cases} 0 \to 0 \to K \to 0 \\ E_1(c) \to K \to E_1(a) \to E_1(c)[1] \end{cases}$$

the first triangle is a weight-$(c-a)$ exact triangle, since it is readily verified that K is $(c-a)$-acyclic. Then $\delta^{\mathcal{F}}(E_1(a), E_1(c)) \leq (c-a) + 0 \leq \tau$ by (iv) in the definition of d_{bot} above. On the other hand, consider the following cone decomposition with linearization $(0, E_1(a), 0)$ (note that, by definition, $\Sigma^{-(a-c)}E_1(a) = E_1(c)$):

$$\begin{cases} 0 \to 0 \to 0 \to 0 \\ E_1(a) \to 0 \to \Sigma^{-(a-c)}E_1(a) \to \Sigma^{-(a-c)}E_1(a)[1] \\ 0 \to E_1(c) \to E_1(c) \to 0, \end{cases} \tag{2.49}$$

where the second triangle has weight $a - c > 0$ by Remark 2.50 (b). Therefore, $\delta^{\mathcal{F}}(E_1(c), E_1(a)) \leq 0 + (a-c) + 0 \leq \tau$, which implies that

$$d^{\mathcal{F}}(E_1(a), E_1(c)) \leq \tau. \tag{2.50}$$

Case II. Estimate $d^{\mathcal{F}}(E_2(a,b), E_2(c,d))$. We will carry out the estimation as follows:

$$d^{\mathcal{F}}(E_2(a,b), E_2(c,d)) \leq d^{\mathcal{F}}(E_2(a,b), E_2(c,b)) + d^{\mathcal{F}}(E_2(c,b), E_2(d,b)).$$

Moreover, we will only estimate $d^{\mathcal{F}}(E_2(a,b), E_2(c,b))$ with $a \geq c$; other situations can be handled in a similar and symmetric way. Similarly to Case I above, consider

the cone decomposition

$$\begin{cases} 0 \to 0 \to K \to 0 \\ E_2(c,b) \to K \to E_2(a,b) \to E_2(c,b)[1], \end{cases}$$

where $E_2(a,b) \to E_2(c,b)$ is the identity map $\langle x \rangle_{E_2(a,b)} \to \langle x \rangle_{E_2(c,b)}$ (and similarly for the generator y), with a negative filtration shift, and K is the cone. Since K is $(a-c)$-acyclic, we have $\delta^{\mathcal{F}}(E_2(a,b), E_2(c,b)) \le 0 + (a-c) \le \tau$. On the other hand,

$$\begin{aligned} &\delta^{\mathcal{F}}(E_2(c,b), E_2(a,b)) \\ &\quad \le \delta^{\mathcal{F}}(E_2(c,b), \Sigma^{a-c} E_2(c,b)) + \delta^{\mathcal{F}}(E_2(a,b+a-c), E_2(a,b)), \end{aligned}$$

where $\delta^{\mathcal{F}}(E_2(c,b), \Sigma^{a-c} E_2(c,b)) \le a-c$ by a similar cone decomposition to that in (2.49). Meanwhile, since $b+a-c \ge b$, the identity map from $E_2(a,b+a-c)$ to $E_2(a,b)$ (with negative filtration shift) yields $\delta^{\mathcal{F}}(E_2(a,b+a-c), E_2(a,b)) \le a-c$. Therefore, together we have, by (iii) in the definition of d_{bot} above,

$$d^{\mathcal{F}}(E_2(a,b), E_2(c,b)) \le 2(a-c) \le 2\tau,$$

which implies

$$d^{\mathcal{F}}(E_2(a,b), E_2(c,d)) \le 4\tau. \tag{2.51}$$

Therefore, by (2.50) and (2.51) together, we have

$$\begin{aligned} d^{\mathcal{F}}(X,Y) &\le \#|\mathcal{B}(X) \setminus \mathcal{B}(X)_{\mathrm{short}}| \cdot 4\tau + \tau \\ &\le (4\,\#|\mathcal{B}(X) \setminus \mathcal{B}(X)_{\mathrm{short}}| + 1)(d_{\mathrm{bot}}(X,Y) + \epsilon) \\ &\le (4 \min\{\#|\mathcal{B}(X)|, \#|\mathcal{B}(Y)|\} + 1)(d_{\mathrm{bot}}(X,Y) + \epsilon), \end{aligned}$$

where the last inequality holds since σ is a bijection by (ii) in the definition of d_{bot} above. Let $\epsilon \to 0$, and we complete the proof. ∎

2.5.3 Topological spaces +

There are many topological categories consisting of topological spaces endowed with additional structures (indicated by the + in the title of the subsection) that can be analyzed with the tools discussed before. We will discuss here two elementary examples. They both fit the following scheme: we will have a triple consisting of a (small) category $\mathcal{K}$, an endofunctor $T_{\mathcal{K}} : \mathcal{K} \to \mathcal{K}$, and a class of triangles $\Delta_{\mathcal{K}}$ in $\mathcal{K}$ of the form

$$A \to B \to C \to T_{\mathcal{K}} A.$$

In these cases, the objects of $\mathcal{K}$ have an underlying structure as topological spaces and, similarly, the morphisms in $\mathcal{K}$ are continuous maps, while the functor $T_{\mathcal{K}}$ corresponds to the suspension of spaces.

The aim is to define fragmentation pseudometrics on the objects of $\mathcal{K}$ by first associating to the triangles in $\Delta_{\mathcal{K}}$ a weight $\overline{w}_{\mathcal{K}} : \Delta_{\mathcal{K}} \to \mathbb{R}$ with some reasonable properties, and then defining the quantities $\delta^{\mathcal{F}}(X, Y)$ and $\underline{\delta}^{\mathcal{F}}(X, Y)$ as in, respectively, (2.2) and (2.5), taking into account only decompositions appealing to triangles $\Delta_i \in \Delta_{\mathcal{K}}$. Notice that $\delta^{\mathcal{F}}$ is not generally defined in this setting as its definition requires to desuspend spaces. On the other hand, as soon as $\overline{w}_{\mathcal{K}}$ is given, $\underline{\delta}^{\mathcal{F}}$ can be defined by formula (2.39) with $w_{\mathcal{K}}$ replacing w_∞ there, and with each triangle in the sequence (2.38) being replaced with a triangle in $\Delta_{\mathcal{K}}$. We assume that the family $\mathcal{F}$ is such that $0 \in \mathcal{F}$, and in most cases we assume implicitly that $\mathcal{F}$ consists of all the objects F such that there are triangles in $\Delta_{\mathcal{K}}$ of the form $F \to A \to B \to T_{\mathcal{K}} F$. The resulting $\underline{\delta}^{\mathcal{F}}$ trivially satisfies the triangle inequality. The pseudometric $\underline{d}^{\mathcal{F}}$ obtained by the symmetrization of $\underline{\delta}^{\mathcal{F}}$ exists in this case too (see Remark 2.5 (c)). Based on the various constructions discussed earlier in the memoir, there are two approaches to define a weight $\overline{w}_{\mathcal{K}}$ (that is not flat), and they both require some more structure:

(A) The additional structure in this case is a functor $\Phi : \mathcal{K} \to \mathcal{C}_\infty$, where $\mathcal{C}$ is a TPC; in the examples below, $\mathcal{C} = H^0 \mathcal{F}\mathcal{K}_{\mathbf{k}}$ – the triangulated persistence homotopy category of filtered cochain complexes. We also require that Φ commutes with T (at least up to some natural equivalence), and that for each $\Delta \in \Delta_{\mathcal{K}}$ the image $\Phi(\Delta)$ of Δ is exact in $\mathcal{C}_\infty$ (and thus $\overline{w}(\Phi(\Delta)) < \infty$, where $\overline{w}$ is the persistence weight introduced in Definition 2.63). In this case, for each $\Delta \in \Delta_{\mathcal{K}}$ we put

$$\overline{w}_{\mathcal{K}}(\Delta) = \overline{w}(\Phi(\Delta)).$$

(B) This second approach requires first that the morphisms $\hom_{\mathcal{K}}(A, B)$ are endowed with a natural increasing filtration compatible with the composition. Secondly, there should be a shift functor $\Sigma_{\mathcal{K}} : (\mathbb{R}, +) \to \mathrm{End}(\mathcal{K})$ compatible with the filtration on morphisms and that commutes with $T_{\mathcal{K}}$. Moreover, the triangles in $\Delta_{\mathcal{K}}$ have to be part of a richer structure, such as a model category or a Waldhausen category (compatible with the functor $\Sigma_{\mathcal{K}}$). In this case, the definition of weighted triangles can be carried out by following the steps in Section 2.3.3, but at the space level, without moving to an algebraic category. This approach goes beyond the scope of this memoir and will not be pursued here.

Remark 2.106. Of course, it is also possible to mix in some sense the two approaches mentioned above. For instance, in the two examples below the category $\mathcal{K}$ carries a shift functor $\Sigma_{\mathcal{K}}$ as in (B), but also a functor Φ as in (A), such that Φ commutes with the shift functors in the domain and target. In that case we can use Φ to pull back

to $\mathcal{K}$ more of the structure and weights in $\mathcal{C}$ (of course, this remains less precise than constructing weights at the space level).

2.5.3.1 Topological spaces with action functionals. We will discuss here a category denoted by $\mathcal{ATop}_*$. The objects of this category are pairs (A, f_A) where $A = (A, *_A)$ is a pointed topological space and $f_A : A \to \mathbb{R}$ is a continuous function bounded from below by the value $f_A(*_A)$ of f_A at the base point $*_A$ of A. We will refer to f_A as the *action functional* associated to A. The morphisms in this category are pointed continuous maps $u : A \to B$ such that there exists $r \in \mathbb{R}$ with the property that $f_B(u(x)) \leq f_A(x) + r$ for all $x \in A$.

We will see that there is a natural *contravariant* functor

$$\Phi : \mathcal{ATop}_* \to [H^0 \mathcal{FK}_{\mathbf{k}}]_\infty \tag{2.52}$$

inducing a weight $\overline{w}_{\mathcal{ATop}_*}$ and the associated pseudometrics $\underline{d}^{\mathcal{F}}$ on $\mathrm{Obj}(\mathcal{ATop}_*)$ along the lines of point (A) above.

Remark 2.107. The condition on f_A being bounded from below is one possible choice in this construction. Its role is to allow the constant map $u : (A, f_A) \to (B, f_B)$ to be part of the morphisms of $\mathcal{ATop}_*$.

Before proceeding with the construction of the functor Φ, we discuss some features of $\mathcal{ATop}_*$. Notice first that the morphisms are filtered, with the rth stage being

$$\begin{aligned}&\hom^{\leq r}_{\mathcal{ATop}_*}(A, B)\\ &\quad = \{u : A \to B \mid u \text{ continuous}, u(*_A) = *_B,\ f_B(u(x)) \leq f_A(x) + r\ \forall x \in A\}.\end{aligned}$$

There is an obvious family of functors $\Sigma_{\mathcal{ATop}_*} : (\mathbb{R}, +) \to \mathcal{ATop}_*$ defined by

$$\Sigma^s_{\mathcal{ATop}_*}(A, f_A) = (A, f_A + s)$$

and being the identity on morphisms. The next step is to define the translation functor $T_{\mathcal{ATop}_*}$. At the underlying topological level this is just the topological suspension, but we need to be more precise about the action functional. Given an object (A, f_A), we first define the cone (CA, f_{CA}). We take CA to be the reduced cone; in other words, the quotient topological space $CA = A \times [0, 1]/(A \times \{1\} \cup *_A \times [0, 1])$. To define f_{CA} we first consider the homotopy $h_A : A \times [0, 1] \to \mathbb{R}$,

$$h_A(x, t) = \begin{cases} f_A(x) & \text{if } 0 \leq t \leq \frac{1}{2}, \\ (2 - 2t)(f_A(x) - f_A(*_A)) + f_A(*_A) & \text{if } \frac{1}{2} \leq t \leq 1. \end{cases}$$

The map $f_{CA} : CA \to \mathbb{R}$ is induced by h_A. We now define the reduced suspension, $SA = CA/A \times \{0\}$, and take f_{SA} to be the map induced to the quotient by the homo-

topy $h'_A : A \times [0,1] \to \mathbb{R}$,

$$h'_A(x,t) = \begin{cases} 2t(f_A(x) - f_A(*_A)) + f_A(*_A) & \text{if } 0 \le t \le \frac{1}{2}, \\ (2-2t)(f_A(x) - f_A(*_A)) + f_A(*_A) & \text{if } \frac{1}{2} \le t \le 1. \end{cases}$$

We put $T(A, f_A) = (SA, f_{SA})$. It is immediate to see that T extends to a functor on $\mathcal{ATop}_*$ and that it commutes with Σ. Moreover, both Σ and T so defined commute and are compatible with the filtration of the morphisms, in the sense that they take $\hom^{\le r}$ to $\hom^{\le r}$ for each r. Moreover, composition of morphisms is also compatible with the filtrations, in the sense that it takes $\hom^{\le r_1}(B,C) \times \hom^{\le r_2}(A,B)$ to $\hom^{\le r_1+r_2}(A,C)$.

We now define the class of exact triangles $\Delta_{\mathcal{ATop}_*}$. For this we consider a morphism $u : (A, f_A) \to (B, f_B)$ and we first define its cone $\mathrm{Cone}(u)$. As a topological space this is, as expected, the quotient topological space $(B \cup CA)/\sim$, where the equivalence relation $\sim$ is generated by $f(x) \sim x \times \{0\}$. The base point of $\mathrm{Cone}(u)$ is the same as that of B. The action functional $f_{\mathrm{Cone}(u)}$ is induced to the respective quotient by

$$G(x) = \begin{cases} f_B(x) & \text{if } x \in B, \\ (1-2t)\big(f_B(u(y)) - f_B(*_B)\big) & \\ \quad + 2t\big(f_A(y) - f_A(*_A)\big) + f_B(*_B) & \text{if } x = (y,t) \in A \times [0, \frac{1}{2}], \\ (2-2t)\big(f_A(y) - f_A(*_A)\big) + f_B(*_B) & \text{if } x = (y,t) \in A \times [\frac{1}{2}, 1]. \end{cases}$$

There is an obvious inclusion $i : (B, f_B) \to \mathrm{Cone}(u)$, and there is also a projection $p : \mathrm{Cone}(u) \to TA$ (that contracts B to a point). This map belongs to our class of morphisms because the functional f_B is bounded from below. The class $\Delta_{\mathcal{ATop}_*}$ consists of triangles Δ:

$$\Delta : A \xrightarrow{u} B \xrightarrow{i} \mathrm{Cone}(u) \xrightarrow{p} TA. \tag{2.53}$$

We finally construct the functor $\Phi : \mathcal{ATop}_* \to [H^0 \mathcal{FK}_{\mathbf{k}}]_\infty$. This functor will be contravariant, since the objects of $\mathcal{FK}_{\mathbf{k}}$ are cochain complexes (rather than chain complexes).

First we fix some notation. For a pointed topological space X, we denote by $\widetilde{C}_*(X)$ the reduced singular chain complex of X with coefficients in $\mathbf{k}$, and by $\widetilde{C}^*(X)$ the reduced singular cochain complex. We denote by $C_*(X)$ and $C^*(X)$ the corresponding non-reduced complexes. If $Y \subset X$ is a pointed subspace, then $\widetilde{C}_*(X,Y)$ and $\widetilde{C}^*(X,Y)$ are the relative (co)chains. Consider an object of $\mathcal{ATop}_*$, (A, f_A), and let $A^{\le r} = (f_A)^{-1}(-\infty, r]$. Notice that the spaces $A^{\le r}$ are pointed (if non-void). There is a filtration of $C^*(A)$ defined by

$$\widetilde{C}^*(A)^{\le -r} = \mathrm{im}\big\{\widetilde{C}^*(A, A^{\le r}) \to \widetilde{C}^*(A)\big\}.$$

Thus the filtration up to $s \in \mathbb{R}$ of $\widetilde{C}^*(A)$ consists of the cochains in A that vanish over the singular chains of $A^{\leq -s} \subset A$. It is clear that the cochain differential preserves this filtration. Moreover, the filtration is increasing, and if $f \in \hom^{\leq r}_{\mathcal{ATop}_*}(A, B)$, then f pulls back the cochains in B that vanish over $B^{\leq a}$ to cochains in A that vanish over $A^{\leq a-r}$. As a result, $C^*(f) : \widetilde{C}^*(B) \to \widetilde{C}^*(A)$ shifts filtration by r. Finally, we define the functor Φ. For each object (A, f_A) of $\mathcal{ATop}_*$ we take $\Phi(A, f_A)$ to consist of the cochain complex $\widetilde{C}^*(A)$ together with the filtration $\{\widetilde{C}^*(A)^{\leq r}\}$ defined above. For a morphism $u : (A, f_A) \to (B, f_B)$ we take $\Phi(u) = [\widetilde{C}^*(u)]$, where $[-]$ represents the cochain-homotopy class of the respective cochain morphism.

The definition of the morphisms in $\mathcal{ATop}_*$ implies that $\Phi(u)$ is indeed a morphism in $[H^0\mathcal{FK}_{\mathbf{k}}]_\infty$. Moreover, because we are using everywhere reduced cochain complexes (and we work in the pointed category), we have that $\Phi(\Delta)$ is exact in $[H^0\mathcal{FK}_{\mathbf{k}}]_\infty$ for each of the triangles in $\Delta_{\mathcal{ATop}_*}$. Further, the functor Φ also interchanges the shift functors in the domain and the target.

In all cases, the weight $\bar{w}_{\mathcal{ATop}_*}$ is well defined, as are the associated fragmentation pseudometrics $\underline{d}^{\mathcal{F}}(-,-)$ on the objects of $\mathcal{ATop}_*$. Roughly speaking, these fragmentation pseudometrics measure how much "weight" is needed to obtain a given topological space via successive cone attachments of spaces in $\mathcal{F}$.

Remark 2.108. (a) The choice of the class $\Delta_{\mathcal{ATop}_*}$ given above is quite restrictive, with the consequence that the resulting pseudometrics are often infinite. One alternative is to enlarge this class to all triangles in $\mathcal{ATop}_*$ that are homotopy equivalent to those in the initial class through maps (and homotopies) of filtration 0.

(b) From some points of view, working in the *pointed* category of spaces endowed with an action functional is not natural. Other choices are possible, in particular ones for which the translation functor T more closely imitates dynamical stabilization.

(c) The restriction of Φ to compact topological spaces admits an obvious lift to $H^0\mathcal{FK}_{\mathbf{k}}$. However, without this restriction, such a lift does not seem to be available in full generality.

2.5.3.2 Metric spaces. The category $\mathcal{M}etr_0$ that we will consider here has as objects path-connected metric spaces (X, d_X) of finite diameter. The morphisms are Lipschitz maps. Recall that $\phi : X \to Y$ is a Lipschitz map if there exists a constant $c \in [0, \infty)$, called the Lipschitz constant of ϕ, with the property that $d_Y(\phi(x), \phi(y)) \leq c\, d_X(x, y)$ for all $x, y \in X$.

Remark 2.109. The finite-diameter condition imposed here – indicated by the subscript 0 – is necessary for some of the constructions below. The connectivity assumption is more a matter of convenience.

We will construct a functor as in (2.52) with one main modification. For convenience, we prefer to define a covariant functor and thus our target category will not

be a category of cochain complexes but rather one of filtered *chain complexes* (the passage from one to the other is formal, replacing C^* by C_{-*} and vice versa). We will denote the category of filtered chain complexes over $\mathbf{k}$ by $\mathcal{FK}'_{\mathbf{k}}$. This behaves just as a usual dg-category, except that the differential on the space of morphisms is of degree -1. With this change, we will construct

$$\Phi' : \mathcal{M}etr_0 \to [H_0 \mathcal{FK}'_{\mathbf{k}}]_\infty \tag{2.54}$$

as well as related structures on $\mathcal{M}etr_0$, as in point (A) at the beginning of the section (see also Remark 2.106).

We start by noting that there is an obvious increasing filtration of the morphisms in $\mathcal{M}etr_0$, with

$$\hom^{\leq r}_{\mathcal{M}etr_0}(X, Y) = \{u : X \to Y \mid \text{the Lipschitz constant of } u \text{ is at most } e^r\}.$$

It is immediate to see that this filtration is compatible with composition. There is also a family of functors $\Sigma_{\mathcal{M}etr_0} : (\mathbb{R}, +) \to \mathcal{M}etr_0$ defined by rescaling the metric, $\Sigma^s(A, d_A) = (A, e^s d_A)$, and by acting as the identity on morphisms. As in the example in the previous section, we will next define the translation functor $T_{\mathcal{M}etr_0}$ and the class of triangles $\Delta_{\mathcal{M}etr_0}$. The first step is to construct the metric cone $C'A$ for an object (A, d_A) in our class. Topologically, the cone $C'A$ will be this time the *unreduced* cone over A. Thus it is defined by $C'A = A \times [0, 1]/A \times \{0\}$. To define the metrics $d_{C'A}$, first let D_A be the diameter of A. We then put

$$d_{C'A}((x, t), (y, t')) = \frac{D_A}{2}|t - t'| + \min\{t, t'\}\, d_A(x, y).$$

It is immediate to see that this does indeed define a metric on $C'A$. A similar construction yields $T(A, d_A)$. Topologically, we will first define the (*non-reduced*) suspension $S'A$ as the topological quotient of $A \times [-\frac{1}{2}, \frac{1}{2}]$, with $A \times \{-\frac{1}{2}\}$ identified to a point S and $A \times \{+\frac{1}{2}\}$ identified to a different point N. We now define $d_{S'A}$ by

$$d_{S'A}((x, t), (y, t')) = \frac{D_A}{2}|t - t'| + \min\{\tfrac{1}{2} - |t|, \tfrac{1}{2} - |t'|\} d_A(x, y),$$

and again it is immediate to see that this defines a metric on $S'A$. We now put $T(A, d_A) = (S'A, d_{S'A})$. The next step is to define the triangles in $\Delta_{\mathcal{M}etr_0}$. For this we assume that $u : (A, d_A) \to (B, d_B)$ is a morphism in our category, and we want to define the (non-reduced) cone of u, $\mathrm{Cone}'(u)$. Topologically, this is, as usual, $B \cup C'A/[\{x\} \times \{0\} \sim u(x) \mid x \in A]$. To define a metric on $\mathrm{Cone}'(u)$, we first notice that, given a map $g : X \to Y$ and a pseudometric d_Y on Y, there is a pullback pseudometric on X given by $g^* d_Y(a, b) = d_Y(g(a), g(b))$. We now let $A' = u(A) \subset B$ and we denote by $\bar{u} : A' \hookrightarrow B$ the inclusion. Notice that $\mathrm{Cone}'(\bar{u}) \subset C'B$. Thus $\mathrm{Cone}'(\bar{u})$ is endowed with a metric given by the restriction of the metric $d_{C'B}$ on $C'B$. There

are obvious projections $\pi : \mathrm{Cone}'(u) \to \mathrm{Cone}'(\bar{u})$ and $p : \mathrm{Cone}'(u) \to S'A$. Here p collapses B to the point S in the suspension and sends each point $(x, t) \in C'A$ to $(x, t - \frac{1}{2})$. We now define

$$d_{\mathrm{Cone}'(u)} := \pi^* d_{C'B} + p^* d_{S'A}.$$

Notice that if u is not injective and B is not a single point, then both pseudometrics in the right-hand side of the equality are degenerate. Nonetheless, $d_{\mathrm{Cone}'(u)}$ is non-degenerate. Finally, the class of triangles $\Delta_{\mathcal{M}etr_0}$ consists of triangles

$$A \xrightarrow{u} B \xrightarrow{i} \mathrm{Cone}'(u) \xrightarrow{p} S'A,$$

where i is the inclusion and p is the projection above.

With this preparation, we can now define the functor Φ' from (2.54). Consider an object (A, d_A) in our category and the associated singular complex $C_*(A)$. This chain complex is filtered as follows:

$$C_k^{\leq r}(A) = \left\{ \sum_i a_i \sigma_i \mid a_i \in \mathbf{k},\ \sigma_i \text{ a singular simplex of diameter at most } e^r \right\}.$$

In other words, in the expression above, $\sigma_i : \Delta^k \to A$ is a continuous map with the standard k-simplex as domain and such that $d_A(\sigma_i(x), \sigma_i(y)) \leq e^r$ for any $x, y \in \Delta^k$. Consider the constant map $c : A \to *$. This induces an obvious surjection $C_* \to C_*(*)$, and we denote by $\overline{C}_*(A)$ the kernel of this map (this is quasi-isomorphic to the reduced singular chain complex of A – because A is connected – but is independent of the choice of base-point). There is an induced filtration $\overline{C}_*^{\leq r}(A)$. We now put

$$\Phi'(A, d_A) = \overline{C}_*(A) \quad \text{with the filtration } \{\overline{C}_*^{\leq r}(A)\}_r.$$

Further, for a morphism $u : (A, d_A) \to (B, d_B)$ we take $\Phi'(u) = [C_*(u)]$, the chain homotopy class of the singular chain map $C_*(u)$ (restricted to the $\overline{C}(-)$ complexes).

It is easy to see that this Φ' is indeed a functor as desired and that $\Phi'(\Delta)$ is exact for each triangle Δ as defined above and, again, Φ' interchanges the shift functors in the domain and target. In summary, the weight $\overline{w}_{\mathcal{M}etr_0}$ is well defined, as are the quantities $\underline{\delta}^{\mathcal{F}}$ and the pseudometrics associated to them.

Remark 2.110. (a) Similarly to Remark 2.108, the definition of the triangles in $\Delta_{\mathcal{M}etr_0}$ is highly restrictive and, in this case, even the objects in our category are subject to a constraint – finiteness of the diameter – that might be a hindrance in applications. One way to apply the methods above to the study of spaces of infinite diameter is to consider triangles of the form $\Delta : A \to B \to C \to S'A$, where A is of finite diameter and such that $S'A$ admits a metric as above, and to analyze when $\Phi'(\Delta)$ is of finite persistence weight in $[H_0 \mathcal{F}\mathcal{K}'_{\mathbf{k}}]_\infty$.

(b) In studying metric spaces of infinite diameter by these methods, it is likely that the most appropriate structure fitting with the cone construction is that of a *length structure*, in the sense of Gromov, as in [35, Chapter 1, Section A]. We will not pursue this theme further here.

2.5.3.3 Further remarks on topological examples. (A) In the topological examples above – for instance in $\mathcal{ATop}_*$ – it is natural to see what the quantities $\underline{\delta}^{\mathcal{F}}(-)$ mean even for the flat weight w_0, which associates to each exact triangle the value 1. Of course, in this case $\underline{\delta}^{\mathcal{F}}(A, B)$ simply counts the minimal number of cone attachments in the category $\mathcal{ATop}_*$ needed to obtain A out of the space B by attaching cones over spaces in the family $\mathcal{F}$ using the family of triangles $\Delta_{\mathcal{ATop}_*}$. Given that the weight is flat, the question is independent of filtrations and shift functors, and it reduces to the identical question in the category of pointed spaces $\mathcal{Top}_*$. In the examples below we will focus on this category and on $\underline{\delta}^{\mathcal{F}}(A, *)$, which is one of the most basic quantities involved.

It is useful to keep in mind that there are two more choices that are essential in defining $\underline{\delta}^{\mathcal{F}}(A, *)$: the choice of family $\mathcal{F}$ and the choice of the class of exact triangles $\Delta_{\mathcal{Top}_*}$; see also Remark 2.108 (a).

(i) $\mathcal{F} = \{S^0, S^1, \ldots, S^k, \ldots\}$; $\Delta_{\mathcal{Top}_*}$ are the triangles $A \xrightarrow{u} B \to \mathrm{Cone}(u) \to SA$ as in (2.53) (but omitting the action functionals). In this case, $\underline{\delta}^{\mathcal{F}}(A, *) = k < \infty$ means that A has the structure of a finite CW-complex with k cells.

(ii) $\mathcal{F} = \{S^0, S^1, \ldots, S^k, \ldots\}$; we now take $\Delta_{\mathcal{Top}_*}$ to be the triangles that are homotopy equivalent to the triangles $A \xrightarrow{u} B \to \mathrm{Cone}(u) \to SA$ from (2.53). In this case, $\underline{\delta}^{\mathcal{F}}(A, *) = k$ means that A is homotopy equivalent to a CW-complex with k cells. This number is obviously a homotopy invariant. It is clearly bounded from below by the sum of the Betti numbers of A.

(iii) $\mathcal{F}$ consists of all pointed spaces with the homotopy type of CW-complexes; $\Delta_{\mathcal{Top}_*}$ are as at (ii). In this case, the definition of $\underline{\delta}^{\mathcal{F}}(A, *)$ coincides with that of the cone-length, $\mathrm{Cl}(A)$, of A (for a space A with the homotopy type of a CW-complex). Cone-length is a homotopical invariant of interest because it is at most one greater than the Lusternik–Schnirelmann category [19], which, in turn, provides a lower bound for the minimal number of critical points of smooth functions on manifolds. Incidentally, as noted by Smale [54], a version of the Lusternik–Schnirelmann category provides also a measure of the complexity of algorithms; see [22] for more on this subject.

(iv) At this point, we will change the underlying category and place ourselves in the pointed category of finite-type, simply-connected *rational* spaces $\mathcal{Top}_1^{\mathbb{Q}}$ (see [31]). We take $\mathcal{F}$ to consist of finite wedges of rational spheres of

dimension at least 2. The triangles $\Delta_{\mathcal{T}op_*^{\mathbb{Q}}}$ are as in (ii) (in the category of rational spaces), but we will also allow in $\Delta_{\mathcal{T}op_*^{\mathbb{Q}}}$ "formal" triangles of the form $S^{-1}F \to * \to F$, where $F \in \mathcal{F}$ (desuspending is not possible in our category, but we still want to have, for a rational 2-sphere $S^2_{\mathbb{Q}}$, $\underline{\delta}^{\mathcal{F}}(S^2_{\mathbb{Q}}, *) = 1$). In this setting, it turns out that

$$\underline{\delta}^{\mathcal{F}}(A, *) = \mathrm{Cl}(A) = \mathrm{nil}(A)$$

(see [20]). Both equalities here are non-trivial: the first because in the definition of Cl we use cones over arbitrary (rational) spaces, while in this example $\mathcal{F}$ consists of only wedges of spheres. For the second equality, $\mathrm{nil}(A)$ is the minimal order of nilpotence of the augmentation ideal $\overline{\mathcal{A}}$ of a rational differential graded commutative algebra $\mathcal{A}$ representing A (recall that, by a celebrated result of Sullivan [55], the homotopy category of rational simply-connected spaces is equivalent to the homotopy category of rational differential graded commutative algebras, the representative of a given space being given by the so-called PL-de Rham complex of A).

(B) One of the difficulties of extracting a triangulated persistence category from a topological category such as those considered in this section is very basic and has to do with the difference between stable and unstable homotopy. In essence, recall that if $\mathcal{C}$ is a TPC, then the 0-level category $\mathcal{C}_0$ is required to be triangulated. However, in unstable settings, homotopy categories of spaces are not triangulated.

(i) An instructive example is a variant of our discussion concerning the category $\mathcal{M}etr_0$. In this case the morphisms $\hom_{\mathcal{M}etr_0}(A, B)$ carry an obvious topology as well as a filtration, as described in Section 2.5.3.2. We can now consider a new category, $\widetilde{\mathcal{M}}etr_0$, with the same objects as $\mathcal{M}etr_0$ but with morphisms $\hom_{\widetilde{\mathcal{M}}etr_0}(A, B) = S_*(\hom_{\mathcal{M}etr_0}(A, B))$, where $S_*(-)$ stands for cubical chain complexes. These morphisms carry an obvious filtration obtained by applying the cubical chains to the filtration of $\hom_{\mathcal{M}etr_0}(A, B)$. The composition in this category is given by applying cubical chains to the composition $\hom_{\mathcal{M}etr_0}(B, C) \times \hom_{\mathcal{M}etr_0}(A, B) \to \hom_{\mathcal{M}etr_0}(A, C)$ and composing with the map

$$\begin{aligned} &S_*(\hom_{\mathcal{M}etr_0}(B, C)) \otimes S_*(\hom_{\mathcal{M}etr_0}(A, B)) \\ &\quad \to S_*(\hom_{\mathcal{M}etr_0}(B, C) \times \hom_{\mathcal{M}etr_0}(A, B)) \end{aligned}$$

induced by taking products of cubes. It follows that $\widetilde{\mathcal{M}}etr_0$ is a filtered dg-category (in homological formalism). Thus all the machinery in Section 2.5.1 is applicable in this case. Moreover, this category carries an obvious shift functor. However, $[H_0\widetilde{\mathcal{M}}etr_0]_\infty$ is not triangulated and thus $\widetilde{\mathcal{M}}etr_0$

is not pre-triangulated (quite far from it). Indeed, $\hom_{[H_0\tilde{\mathcal{M}}etr_0]_\infty}(A, B)$ is the free abelian group generated by the homotopy classes of Lipschitz maps from A to B. As a result, the translation functor (which is in our case the topological suspension) is certainly not an isomorphism.

(ii) As mentioned before, in point (B) at the beginning of Section 2.5.3, one way to bypass these issues is to introduce a sort of filtered Waldhausen category or a similar formalism and to develop a machinery parallel to that of TPCs in this unstable context. The structure present in $\mathcal{ATop}_*$ and $\mathcal{M}etr_0$ suggests that such a construction is possible and will be relevant in these cases.

(iii) There is yet another, more geometric, approach to associating to each of $\mathcal{ATop}_*$ and $\mathcal{M}etr_0$ a triangulated persistence category. This is based on moving from these categories to stable categories, where the underlying objects are the spectra obtained by stabilizing the objects of the original categories and the morphisms come with an appropriate filtration induced from the respective structures (action functionals or, respectively, metrics) on the initial objects. This seems likely to work and to directly produce a TPC, but we will not pursue the details at this time.

2.5.4 Filtrations in Tamarkin's category

This section is devoted to an example of a triangulated persistence category that comes from the filtration structure present in Tamarkin's category. This category was originally defined in [56], based on singular supports of sheaves, and was used to prove some non-displaceability results in symplectic geometry, as well as other more recent results related to Hamiltonian dynamics (see [36]).

2.5.4.1 Background on Tamarkin's category. Let X be a manifold, and let $\mathcal{D}(\mathbf{k}_X)$ be the derived category of sheaves of $\mathbf{k}$-modules over X. In particular, this is a triangulated category. For any $A \in \mathrm{Obj}(\mathcal{D}(\mathbf{k}_X))$, due to microlocal sheaf theory as established in [40], one can define the singular support of A, denoted by $\mathrm{SS}(A)$, as a conical (singular) subset of T^*X. We refer to [40, Chapter V] for the precise definition of $\mathrm{SS}(A)$ and a detailed study of its properties. Now, let $X = M \times \mathbb{R}$, where M is a closed manifold, and denote by τ the co-vector coordinate of $T^*\mathbb{R}$ in $T^*(M \times \mathbb{R})$. Consider the full subcategory $\mathcal{D}_{\{\tau \leq 0\}}(\mathbf{k}_{M\times\mathbb{R}})$ of $\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})$ defined by

$$\mathrm{Obj}(\mathcal{D}_{\{\tau\leq 0\}}(\mathbf{k}_{M\times\mathbb{R}})) = \{A \in \mathrm{Obj}(\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})) \mid \mathrm{SS}(A) \subset \{\tau \leq 0\}\}.$$

If $A \to B \to C \to A[1]$ is an exact triangle in $\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})$, then $\mathrm{SS}(C) \subset \mathrm{SS}(A) \cup \mathrm{SS}(B)$. This implies that $\mathcal{D}_{\{\tau\leq 0\}}(\mathbf{k}_{M\times\mathbb{R}})$ is a triangulated subcategory of $\mathcal{D}_{\mathbf{k}_{M\times\mathbb{R}}}$.

Tamarkin's category is defined by

$$\mathcal{T}(M) := \mathcal{D}_{\{\tau\leq 0\}}(\mathbf{k}_{M\times\mathbb{R}})^{\perp,l}, \tag{2.55}$$

where $\perp, l$ denotes the left orthogonal complement of $\mathcal{D}_{\{\tau\leq 0\}}(\mathbf{k}_{M\times\mathbb{R}})$ in $\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})$. Then $\mathcal{T}(M)$ is also a triangulated subcategory. Note that $\mathcal{T}(M) \subset \mathcal{D}_{\{\tau\geq 0\}}(\mathbf{k}_{M\times\mathbb{R}})$ by definition. When $M = \{\text{pt}\}$, Tamarkin's category $\mathcal{T}(\{\text{pt}\})$, together with a constructibility condition, can be identified with the category of persistence $\mathbf{k}$-modules (see [63, Section A.1]).

Remark 2.111. There exists a restricted version of Tamarkin's category, denoted by $\mathcal{T}_V(M)$, where $V \subset T^*M$ is a closed subset (see [63, Section 3.2]). This restricted Tamarkin category is useful to prove the non-displaceability of some subsets in T^*M (see [2]). In this memoir, we will only focus on $\mathcal{T}(M)$.

One way to understand the definition (2.55) is that $\mathcal{T}(M)$ is an admissible subcategory (see [44, Definition 1.8]), in the sense that for any object A in $\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})$ one can always split A in the form of an exact triangle

$$B \to A \to C \to B[1] \tag{2.56}$$

in $\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})$, where $B \in \mathcal{T}(M)$ and $C \in \mathcal{D}_{\{\tau\leq 0\}}(\mathbf{k}_{M\times\mathbb{R}})$. In fact, this splitting can be achieved in a rather concrete manner, which involves an important operator called sheaf convolution on objects in $\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})$. Explicitly, for any two objects A, B in $\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})$, the sheaf convolution of A and B is defined by

$$A * B := \delta^{-1} Rs_!(\pi_1^{-1}A \otimes \pi_2^{-1}B),$$

where $\pi_i : (M\times\mathbb{R})^2 \to M\times\mathbb{R}$ are the projections to each factor of $M\times\mathbb{R}$, s leaves the $(M\times M)$-part unchanged but adds up two inputs on the $\mathbb{R}$-factors, and δ is the diagonal embedding from M to $M\times M$. For instance, $\mathbf{k}_{M\times[0,\infty)} * \mathbf{k}_{M\times[0,\infty)} = \mathbf{k}_{M\times[0,\infty)}$, where $\mathbf{k}_V$ for a closed subset V denotes the constant sheaf supported on V. Moreover, this operator is commutative and associative. An important characterization of an object in $\mathcal{T}(M)$ is that (see [56, Proposition 2.1])

$$A \in \text{Obj}(\mathcal{T}(M)) \quad \text{if and only if} \quad A * \mathbf{k}_{M\times[0,\infty)} = A,$$

which implies that (i) for any object A in $\mathcal{D}(\mathbf{k}_{M\times\mathbb{R}})$, the sheaf convolutions $B := A * \mathbf{k}_{M\times[0,\infty)}$ and $C := A * \mathbf{k}_{M\times(0,\infty)}[1]$ provide the desired exact triangle for a splitting of A in (2.56); (ii) sheaf convolution is a well-defined operator on $\mathcal{T}(M)$.

With the help of sheaf convolution, the $\mathbb{R}$-component generates a filtration structure in $\mathcal{T}(M)$ as follows. For any $r \in \mathbb{R}$, consider the map $T_r : M\times\mathbb{R} \to M\times\mathbb{R}$ defined by $(m, a) \mapsto (m, a + r)$. One can show that for any object A in $\mathcal{T}(M)$, the induced object is $(T_r)_*A = A * \mathbf{k}_{M\times[r,\infty)}$ (see [63, Lemma 3.2]). In fact, $\{(T_r)_*\}_{r\in\mathbb{R}}$

defines an $\mathbb{R}$-family of functors on $\mathcal{T}(M)$. Moreover, if $r \leq s$, then by the restriction map $\mathbf{k}_{M\times[r,\infty)} \to \mathbf{k}_{M\times[s,\infty)}$, we have a canonical morphism $\tau_{r,s}(A) : (T_r)_* A \to (T_s)_* A$. At this point, notice that for $r \leq s$ there does not exist a non-zero morphism from $\mathbf{k}_{M\times[s,\infty)}$ to $\mathbf{k}_{M\times[r,\infty)}$, so the canonical map $\tau_{r,s}$ respects the partial order $\leq$ on $\mathbb{R}$. For any $r \leq s$, $\tau_{r,s}$ is viewed as a natural transformation from $(T_r)_*$ to $(T_s)_*$. Finally, we call an object A in $\mathcal{T}(M)$ a c-torsion element if $\tau_{0,c}(A) : A \to (T_c)_* A$ is zero. For instance, when $M = \{\mathrm{pt}\}$, the constant sheaf $\mathbf{k}_{[a,b)} \in \mathcal{T}(\{\mathrm{pt}\})$ with a finite interval $[a, b)$ is a $(b - a)$-torsion.

We end this subsection with a discussion of the hom-set in $\mathcal{T}(M)$. It is more convenient to consider derived hom, that is, $\mathrm{Rhom}_{\mathcal{T}(M)}(A, B)$ for any two objects A, B in $\mathcal{T}(M)$. Lemma 3.3 in [63] (or Lemma 3.8 (1) in [56]) provides a more explicit way to express such Rhom, that is,

$$\mathrm{Rhom}_{\mathcal{T}(M)}(A, B) = \mathrm{Rhom}_{\mathbb{R}}(\mathbf{k}_{[0,\infty)}, R\pi_* \mathcal{H}om^*(A, B)). \tag{2.57}$$

By taking the cohomology at degree 0, we obtain $\hom_{\mathcal{T}(M)}(A, B)$ as a $\mathbf{k}$-module. Here, $\pi : M \times \mathbb{R} \to \mathbb{R}$ and $\mathcal{H}om^*(\cdot, \cdot)$ is the right adjoint functor to the sheaf convolution (see [2, Definition 3.1]). The right-hand side of (2.57) is relatively computable, since they are all (complexes of) sheaves over $\mathbb{R}$ (cf. [63, Section A.2]). Moreover, by using the adjoint relation between $\mathcal{H}om^*(\cdot, \cdot)$ and the sheaf convolution, one obtains a shifted version of (2.57), that is,

$$\mathrm{Rhom}_{\mathcal{T}(M)}(A, (T_r)_* B) = \mathrm{Rhom}_{\mathbb{R}}(\mathbf{k}_{[0,\infty)}, (T_r)_*(R\pi_* \mathcal{H}om^*(A, B))). \tag{2.58}$$

Therefore, for any $r \leq s$, there exists a well-defined morphism

$$\iota_{r,s}^{A,B} : \hom_{\mathcal{T}(M)}(A, (T_r)_* B) \to \hom_{\mathcal{T}(M)}(A, (T_s)_* B), \tag{2.59}$$

which is induced by the morphism $\tau_{r,s}(R\pi_* \mathcal{H}om^*(A, B))$. Finally, we have a canonical isomorphism,

$$T_r : \hom_{\mathcal{T}(M)}((T_r)_* A, (T_r)_* B) \simeq \hom_{\mathcal{T}(M)}(A, B), \tag{2.60}$$

which is induced by the sheaf convolution with $\mathbf{k}_{M\times[r,\infty)}$. In particular, T_r commutes with the morphism $\iota_{r,s}^{A,B}$ defined in (2.59).

2.5.4.2 Persistence category from Tamarkin's shift functors. We have seen before that Tamarkin's category is endowed with a shift functor. We now discuss the persistence structure induced by this shift functor; see Remark 2.23 (d).

Definition 2.112. Given the category $\mathcal{T}(M)$ as above, define an enriched category $\mathcal{P}(M)$ as follows. The object set of $\mathrm{Obj}(\mathcal{P}(M))$ is the same as $\mathrm{Obj}(\mathcal{T}(M))$, and the hom-set is defined by

$$\hom_{\mathcal{P}(M)}(A, B) = \left\{\{\hom_{\mathcal{T}(M)}(A, (T_r)_* B)\}_{r\in\mathbb{R}}, \{\iota_{r,s}^{A,B}\}_{r,s\in\mathbb{R}, r\leq s}\right\}$$

for any two objects A, B in $\mathcal{P}(M)$, where $\iota_{r,s}^{A,B}$ is the morphism defined in (2.59).

Remark 2.113. Definition 2.112 can be regarded as a generalization of (2.15) in Section 2.2.4, since when $M = \{\text{pt}\}$, Tamarkin's category $\mathcal{T}(\{\text{pt}\})$ can be identified with the category of persistence $\mathbf{k}$-modules. Also, Definition 2.112 fits with geometric examples. Indeed, recall a concrete computation of $\hom_{\mathcal{T}(M)}(A, B)$ when both A and B are sheaves coming from generating functions on M (see [63, Section 3.9]). In this case, $\hom_{\mathcal{P}(M)}(A, B)$ can be identified to a (Morse) persistence $\mathbf{k}$-module in the classical sense.

Lemma 2.114. *The category $\mathcal{P}(M)$ from Definition 2.112 is a persistence category.*

Proof. Consider the functor $E_{A,B} : (\mathbb{R}, \leq) \to \text{Vect}_{\mathbf{k}}$ defined by

$$E_{A,B}(r)(= \hom^r(A, B)) = \hom_{\mathcal{T}(M)}(A, (T_r)_* B),$$

and, for the morphism $i_{r,s}$ when $r \leq s$, set $E_{A,B}(i_{r,s}) = \iota_{r,s}^{A,B}$. Notice that the composition $E_{A,B}(r) \times E_{B,C}(s) \to E_{A,C}(r+s)$ is well defined due to (2.60). Indeed, for any $f \in E_{A,B}(r)$ and $g \in E_{B,C}(s)$, the composition is defined by

$$(f, g) \mapsto T_r(g) \circ f \in \hom_{\mathcal{T}(M)}(A, (T_{r+s})_* C).$$

Then, for any $r \leq r'$, $s \leq s'$, we have

$$\left(T_{r'}\left(\iota_{s,s'}^{B,C}(g)\right)\right) \circ \iota_{r,r'}^{A,B}(f) = \iota_{r+s,r'+s'}^{A,C}(T_r(g) \circ f),$$

which completes the proof that $\mathcal{P}(M)$ is a persistence category. ∎

We now list some properties of the persistence category $\mathcal{P}(M)$.

(a) The 0-level category $\mathcal{P}(M)_0$ has the same objects as $\mathcal{P}(M)$, but

$$\hom_{\mathcal{P}(M)_0}(A, B) = \hom_{\mathcal{T}(M)}(A, B).$$

We use the fact that $(T_0)_* = \mathbb{1}$. Thus, $\mathcal{P}(M)_0 = \mathcal{T}(M)$. This category is triangulated, as we have seen above.

(b) The ∞-level, $\mathcal{P}(M)_\infty$, has the same objects as $\mathcal{P}(M)$, but

$$\hom_{\mathcal{P}(M)_\infty}(A, B) = \varinjlim_{r \to \infty} \hom_{\mathcal{T}(M)}(A, (T_r)_* B),$$

where the direct limit is taken via the map $\iota_{r,s}^{A,B}$. This limit category has been considered in [37, formula (81) and Proposition 6.7], where it is viewed from the perspective of a categorical localization on torsion elements. This can be regarded as a special case of Proposition 2.39 in Section 2.3.2, where the localization is established for a general triangulated persistence category.

(c) On $\mathcal{P}(M)$, each $(T_r)_*$ is a persistence functor for any $r \in \mathbb{R}$, i.e., $(T_r)_* \in \mathcal{P}(\text{End}(\mathcal{P}(M)))$, since T_r commutes with $\tau_{r,s}$.

(d) There exists a natural shift functor on $\mathcal{P}(M)$. Define

$$\Sigma : (\mathbb{R}, +) \to \mathcal{P}(\mathrm{End}(\mathcal{P}(M)))$$

by $\Sigma(r)(=\Sigma^r) = (T_{-r})_*$. For any $r, s \in \mathbb{R}$ and $\eta_{r,s} \in \hom_{\mathbb{R}}(r, s)$, define

$$\Sigma(\eta_{r,s}) := \mathbb{1}_{(T_{-r})_*(-)}.$$

Then, for any object A in $\mathcal{P}(M)$,

$$\begin{aligned}\Sigma(\eta_{r,s})_A = \mathbb{1}_{(T_{-r})_*A} &\in \hom_{\mathcal{T}(M)}((T_{-r})_*A, (T_{-r})_*A)\\ &= \hom_{\mathcal{T}(M)}((T_{-r})_*A, (T_{s-r})_*((T_{-s})_*A))\\ &= E_{(T_{-r})_*A,(T_{-s})_*A}(s-r)\\ &= \hom^{s-r}((T_{-r})_*A, (T_{-s})_*A).\end{aligned}$$

In other words, $\Sigma(\eta_{r,s})_A$ is a natural transformation of shift $s - r$. In particular, the morphism $\eta_r^A = i_{-r,0}((\eta_{r,0})_A) \in \hom^0(\Sigma^r A, A)$ is well defined for any $r \geq 0$. It is easy to check that $\eta_r^A = (\tau_{-r,0})(A)$.

(e) The r-acyclic objects in $\mathcal{P}(M)$ are precisely the r-torsion elements in $\mathcal{T}(M)$. Indeed, by definition, an object A in $\mathcal{P}(M)$ is r-acyclic if and only if $\eta_r^A = \tau_{-r,0}(A) : (T_{-r})_*A \to A$ is the zero morphism, which coincides with the definition of an r-torsion element under the isomorphism (2.60).

(f) Recall that for each r, $\hom^r(-, X) = \hom_{\mathcal{T}(M)}(-, (T_r)_*X)$. This is an exact functor due to (2.58) on $\mathcal{P}(M)_0 = \mathcal{T}(M)$. Similarly, $\hom^r(X, -)$ is also an exact functor on $\mathcal{P}(M)_0 = \mathcal{T}(M)$.

Lemma 2.115. *For any $r \geq 0$ and any object A in $\mathcal{P}(M)$, the morphism*

$$\eta_r^A : (T_{-r})_*A \to A$$

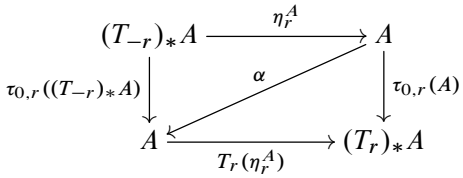
embeds into the exact triangle

$$(T_{-r})_*A \xrightarrow{\eta_r^A} A \to K \to A \tag{2.61}$$

in $\mathcal{T}(M) = \mathcal{P}(M)_0$, where K is r-acyclic.

Proof. Since $\mathcal{T}(M)$ is a triangulated category, the morphism η_r^A embeds into an exact triangle as (2.61). By item (e) above, we need to show that K is an r-torsion element. By [37, Lemma 6.3 (ii)], which provides a criterion for testing whether an object in an exact triangle is a torsion element, it suffices to verify that the following diagram

is commutative for some morphism α:



Indeed, this diagram is commutative if we choose $\alpha = \mathbb{1}_A$ and use the functorial properties of T_{-r}. ∎

Remark 2.116. By the definition of an r-isomorphism given earlier, Lemma 2.115 implies that the morphism $\eta_r^A \in \hom^0((T_r)_*A, A)$ is an r-isomorphism. On the other hand, [63, Section 3.10] defines an interleaving relation between two objects in $\mathcal{T}(M)$, which is similar to r-isomorphism in the sense that A and $(T_r)_*A$ are r-interleaved.

Example 2.117. Let $M = \{\text{pt}\}$ and consider $\mathcal{T}(\{\text{pt}\})$. For $A = \mathbf{k}_{[0,\infty)}$, we know that $(T_{-r})_*A = \mathbf{k}_{[-r,\infty)}$ for any $r \geq 0$. Then we have an exact triangle in $\mathcal{T}(\{\text{pt}\})$:

$$\mathbf{k}_{[-r,\infty)} \xrightarrow{\tau_{-r,0}(A)} \mathbf{k}_{[0,\infty)} \to \mathbf{k}_{[-r,0)}[1] \to \mathbf{k}_{[-r,\infty)}[1],$$

where, as we have seen, $\mathbf{k}_{[-r,0)}[1]$ is an r-torsion element (so r-acyclic). Here, by definition, $\tau_{-r,0}(A)$ is the restriction map from $\mathbf{k}_{[-r,\infty)}$ to $\mathbf{k}_{[0,\infty)}$, and the exact triangle is from [40, (2.6.33)].

The properties stated in (a) and (d) above, together with Lemmas 2.114 and 2.115, imply the main consequence of this section.

Corollary 2.118. *The category $\mathcal{P}(M)$, introduced in Definition 2.112, is a triangulated persistence category.*

Chapter 3

Triangulated persistence Fukaya categories

In this chapter we apply the theory developed in Chapter 2 to the case of Fukaya categories. The setup described before applies naturally to this context: under (significant) constraints, the derived Fukaya category naturally admits a TPC refinement, and this setting is ideal for approaching a variety of quantitative questions typical of symplectic topology.

We begin in Section 3.1 with the statements of the main symplectic applications in the memoir: Theorems 3.1 and 3.4 and Corollary 3.7. To prove these statements we first fix in Section 3.2 the basics of filtered A_∞-categories and associated TPCs, and we then discuss basic notions relative to filtered Floer theory. We describe how to proceed from Floer chain complexes to the Fukaya category. However, for technical reasons the construction leads only to a *weakly* filtered A_∞-category. In Section 3.3 we show that under certain restrictive conditions this construction can be adjusted to obtain a genuinely filtered A_∞-category. The main technical result of Chapter 3 appears in Theorem 3.12. The model for the Fukaya category that we construct in this case is based on clusters of punctured disks. While similar models have appeared before in the literature, we include enough details to justify the control of filtrations. In Section 3.4 we prove the statements from Section 3.1. In particular, we construct the metrics on the spaces of Lagrangians that were announced in the introduction of the memoir. The TPC formalism was inspired by earlier work on Lagrangian cobordism and it is useful to see how weighted triangles and operations with them appear geometrically in the cobordism setting. This is discussed in Section 3.5 together with some other geometric illustrations of some of the statements in Section 3.1.

3.1 Main symplectic topology applications

Let $(X, \omega = d\lambda)$ be a Liouville manifold, i.e., an exact symplectic manifold with a prescribed primitive λ of the symplectic structure ω and such that X is symplectically convex at infinity with respect to these structures. We will work here with pairs $L = (\bar{L}, h_L)$ consisting of a closed oriented exact Lagrangian submanifold $\bar{L} \subset X$ equipped with a function $h_L : \bar{L} \to \mathbb{R}$ that is a primitive of $\lambda|_{\bar{L}}$, i.e., $dh_L = \lambda|_{\bar{L}}$. We will refer to such a pair L as a marked Lagrangian submanifold and to $\bar{L}$ as its underlying Lagrangian.

Fix a collection of marked Lagrangians $\mathcal{X}$ in X. We assume that $\mathcal{X}$ is closed under shifts of the primitives, that is, if $L = (\bar{L}, h_L)$ is in $\mathcal{X}$, then for every $r \in \mathbb{R}$ and $k \in \mathbb{Z}$, the marked Lagrangian $\Sigma^r L := (\bar{L}, h_L + r)$ is also in $\mathcal{X}$. We will also

assume that our marked Lagrangians are graded, in a sense recalled in Section 3.2.2.2. If we need to make the grading explicit, we write $L = (\bar{L}, h_L, \theta_L)$, and we assume the family $\mathcal{X}$ is also closed under translations $L[k] = (\bar{L}, h_L, \theta_L - k)$ of the grading.

Denote by $\bar{\mathcal{X}} = \{\bar{L} \mid L \in \mathcal{X}\}$ the collection of underlying Lagrangian submanifolds corresponding to the marked Lagrangians in $\mathcal{X}$. We will assume that the family $\bar{\mathcal{X}}$ is finite and that its elements are in general position, in the sense that any two distinct Lagrangians $L', L'' \in \bar{\mathcal{X}}$ intersect transversely, and for every three distinct Lagrangians $\bar{L}_0, \bar{L}_1, \bar{L}_2 \in \bar{\mathcal{X}}$ we have $\bar{L}_0 \cap \bar{L}_1 \cap \bar{L}_2 = \emptyset$.

As earlier in the memoir, algebraic considerations can be carried out over an arbitrary field $\mathbf{k}$. However, without additional assumptions on our Lagrangians, Floer theory works only over $\mathbf{k} = \mathbb{Z}_2$. We will therefore assume $\mathbf{k} = \mathbb{Z}_2$, but continue to denote the base field by $\mathbf{k}$, to indicate that under additional assumptions, our theory is expected to work over an arbitrary field $\mathbf{k}$. The marked Lagrangians in X are the objects of an A_∞-category, the Fukaya category $\mathcal{F}uk(X)$ of X, constructed as in Seidel's book [53]. The associated derived Fukaya category is denoted by $D\mathcal{F}uk(X)$. Its objects are the A_∞-modules over $\mathcal{F}uk(X)$ that belong to the triangulated completion of the Yoneda A_∞-modules, $\mathcal{Y}(L)$, where L is a marked Lagrangian. We denote by $\mathcal{F}uk(\mathcal{X})$ the A_∞-subcategory of $\mathcal{F}uk(X)$ with objects the Lagrangians in $\mathcal{X}$, and by $j_{\mathcal{X}} : \mathcal{F}uk(\mathcal{X}) \to \mathcal{F}uk(X)$ the inclusion. There are two Yoneda-type modules associated to the elements of $\mathcal{X}$: over the category $\mathcal{F}uk(X)$ and over the smaller category $\mathcal{F}uk(\mathcal{X})$. The two are related by applying the pullback $j_{\mathcal{X}}^*(-)$ and thus will be generally denoted by the same symbol.

We denote by $D\mathcal{F}uk(\mathcal{X})$ the associated derived category, consisting this time of modules over $\mathcal{F}uk(\mathcal{X})$ that belong to the triangulated completion of the Yoneda modules of the elements of $\mathcal{X}$. We emphasize that, with the terminology used in this memoir, a family $\mathcal{Z}$ of objects in a triangulated category $\mathcal{C}$ is a system of generators of $\mathcal{C}$ if the triangulated envelope of $\mathcal{Z}$ in $\mathcal{C}$ equals $\mathcal{C}$. In particular, the Yoneda modules of the elements of $\mathcal{X}$ form a system of generators of $D\mathcal{F}uk(\mathcal{X})$.

The following consequence of Theorem 3.12 is significant enough to be stated separately:

Theorem 3.1. *There exists a triangulated persistence category $\mathcal{C}\mathcal{F}uk(\mathcal{X})$, independent up to TPC equivalence of the data used in its construction, such that:*

(i) *For any $L, L' \in \mathcal{X}$, there is a canonical isomorphism*

$$\hom_{\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty}(L, L') \cong \mathrm{HF}^0(L, L'),$$

where $\mathrm{HF}^0(-,-)$ is Floer homology in cohomological degree 0.

(ii) *$\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty$ is triangulated equivalent to $D\mathcal{F}uk(\mathcal{X})$.*

(iii) *If the family $\mathcal{X}$ generates $D\mathcal{F}uk(X)$, then for each marked Lagrangian N that intersects transversely the family $\mathcal{X}$, the pullback $j_{\mathcal{X}}^*\mathcal{Y}(N)$ of $\mathcal{Y}(N)$*

– the Yoneda module of N over $\mathcal{F}uk(X)$ (defined with a convenient choice of perturbation data) – is quasi-isomorphic to an object in $\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty$.

We call $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ the *triangulated persistence Fukaya category* associated to $\mathcal{X}$.

Here we emphasize that the construction of $\mathcal{F}uk(\mathcal{X})$ always depends on some perturbation data $\mathcal{P}$; more precisely, $\mathcal{F}uk(\mathcal{X})$ is denoted by $\mathcal{F}uk(\mathcal{X};\mathcal{P})$. Theorem 3.12 guarantees that there always exists perturbation data $\mathcal{P}$ such that $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ is a strict unital A_∞-category, together with filtered A_∞-functors $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0} : \mathcal{F}uk(\mathcal{X};\mathcal{P}_0) \to \mathcal{F}uk(\mathcal{X};\mathcal{P}_1)$ when changing the perturbation data from $\mathcal{P}_0$ to $\mathcal{P}_1$. An essential part of the proof of Theorem 3.1 (see Section 3.4.1.1) shows that the resulting triangulated persistence Fukaya category $\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P})$, constructed from $\mathcal{F}uk(\mathcal{X};\mathcal{P})$, is in fact independent of the perturbation data up to a TPC equivalence. This justifies the notation $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ above without any reference to $\mathcal{P}$.

Point (ii) of Theorem 3.1 implies that if $\mathcal{X}$ generates $D\mathcal{F}uk(X)$, then $\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty$ is equivalent to $D\mathcal{F}uk(X)$. Point (iii) of Theorem 3.1 gives a bit more information and shows that one can use measurements in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ to study Lagrangians that do not necessarily belong to the finite family $\overline{\mathcal{X}}$. Nonetheless, it remains that the requirement that the family $\overline{\mathcal{X}}$ be finite is highly constraining. It is expected that this requirement can be dropped by using a more involved construction in place of the one used in the proof of Theorem 3.12.

Remark 3.2. As stated, Theorem 3.1 identifies $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ up to TPC equivalence (see Definition 2.25) but, while this equivalence is expected to be canonical, our methods do not quite give that. Still, the equivalences that appear here are not completely arbitrary. For example, their mapping on objects leaves the elements of $\mathcal{X}$ fixed. See Theorem 3.12 for more details.

The next result in this section will be formulated in terms of this TPC, $\mathcal{C}\mathcal{F}uk(\mathcal{X})$, and will involve a notion of relative Gromov width that first appeared in [4] (see also [10]). Assume that L and L' are two Lagrangians, both possibly immersed. We define

$$\delta(L;L') = \sup\big\{\pi r^2 \mid \exists\, e : (B(r),\omega_0) \to (X,\omega) \text{ symplectic embedding,} \\ \text{such that } e^{-1}(L) = \mathbb{R}B(r),\ e(B(r)) \cap L' = \emptyset\big\}, \tag{3.1}$$

where $(B(r),\omega_0)$ is the standard closed ball of radius r in $(\mathbb{R}^{2n},\omega_0)$, and $\mathbb{R}B(r) = (\mathbb{R}^n \times \{0\}) \cap B(r)$ is its real part. A related measurement reflects the "quality" of the intersection points between L and L', relative to another subset. Assume that

L and L' intersect transversely and let $A \subset X$ be a subset. We define

$$\begin{aligned}\delta^{\cap}(L, L'; A) = \sup\big\{\pi r^2 \mid\ & \forall x \in L \cap L',\ \exists\, e_x : (B(r), \omega_0) \to (X, \omega) \text{ symp. emb.} \\ & \text{s.t. } e_x(0) = x,\ e_x^{-1}(L) = \mathbb{R}B(r),\ e_x^{-1}(L') = i\mathbb{R}B(r), \\ & e_{x'}(B(r)) \cap e_{x''}(B(r)) = \emptyset \text{ whenever } x' \neq x'', \\ & \text{and moreover } \forall x \in L \cap L',\ e_x(B(r)) \cap A = \emptyset\big\}.\end{aligned} \tag{3.2}$$

Here $i\mathbb{R}B(r) := (\{0\} \times \mathbb{R}^n) \cap B(r)$ is the "imaginary" part of the ball $B(r)$.

We will also need the spectral distance between two marked Lagrangians L and L'. We assume that L is Hamiltonian isotopic to L'. In this case, the Floer homology $\mathrm{HF}(L, L')$ is isomorphic to the singular homology $H_*(L; \mathbf{k})$ of L and there is a canonical class $o_{L,L'} \in \mathrm{HF}(L, L')$ corresponding to the fundamental class in $H_*(L; \mathbf{k})$. There is also a second class $pt_{L,L'} \in \mathrm{HF}(L, L')$ that corresponds to the point class in $H_*(L; \mathbf{k})$. Assume further that $L, L' \in \mathcal{X}$. In this case, given point (i) of Theorem 3.1, we have

$$\mathrm{HF}^0(L, L') = \big\{[f] \in \hom_{\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty}(L, L') \mid f \in \hom_{\mathcal{C}\mathcal{F}uk(\mathcal{X})}(L, L')\big\}.$$

Therefore these classes in $\mathrm{HF}(L, L')$ have spectral numbers $\sigma(-)$ as defined in (2.36). We define

$$\sigma(L, L') = \sigma(o_{L,L'}) - \sigma(pt_{L,L'}).$$

We extend the definition of σ to the case when L' is not Hamiltonian isotopic to L by setting $\sigma(L, L') = \infty$ in this case.

Remark 3.3. It is easily seen that this definition coincides with previous versions of spectral invariants introduced by Viterbo, Schwarz, and Oh, and later adjusted to the Lagrangian setting.

Pick a family $\mathcal{F} \subset \mathcal{X}$ that is invariant with respect to shift and translation. Fix an admissible perturbation data $\mathcal{P}$ and an associated triangulated persistence category $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P})$. Consider the shift-invariant, persistence, fragmentation pseudometric $\widehat{d^{\mathcal{F}}}(-,-)$ associated to the persistence weight $\bar{w}$ on $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P})_\infty$, as described in Section 2.4.3.1 and (2.41). Each such pseudometric is defined on the objects of $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P})$, which contain the Yoneda modules of the Lagrangians in $\mathcal{X}$ but also additional A_∞-modules.

We now define a pseudometric on $\mathcal{X}$ by

$$D^{\mathcal{F}}(L, L') = \widehat{d^{\mathcal{F}}}(\mathcal{Y}(L), \mathcal{Y}(L')), \quad L, L' \in \mathcal{X}.$$

In the case $\mathcal{F} = \{0\}$, we write $D(-,-) = D^{\{0\}}(-,-)$. This is an upper bound for all the other fragmentation metrics $D^{\mathcal{F}}$. There is a slight abuse in notation here because the definition of $D^{\mathcal{F}}$ depends implicitly on the perturbation data $\mathcal{P}$, but this will be resolved in the next result.

Theorem 3.4. *Let $\mathcal{F} \subset \mathcal{X}$. In the setting above, the pseudometrics $D^{\mathcal{F}}$ are independent of the perturbation data $\mathcal{P}$ used for their definitions. Moreover:*

(i) *(Spectrality) Assume that the Lagrangians in $\mathcal{X}$ are graded; then for any $L, L' \in \mathcal{X}$ we have*

$$D(L, L') \leq 4\, \sigma(L, L').$$

(ii) *(Non-degeneracy) For all $L, L' \in \mathcal{X}$,*

$$\frac{\delta(L; L' \cup \bigcup_{F \in \mathcal{F}} F)}{8} \leq D^{\mathcal{F}}(L, L').$$

(iii) *(Persistence of intersections) Assume that $L, L', N \in \mathcal{X}$, $L' \notin \mathcal{F}$. If*

$$D^{\mathcal{F}}(L, L') < \frac{1}{16}\, \delta^{\cap}\big(N, L'; \bigcup_{F \in \mathcal{F}} F\big),$$

then

$$\#(L \cap N) \geq \#(L' \cap N).$$

(iv) *(Finiteness) If the family $\mathcal{F}$ generates $D\mathcal{F}uk(X)$, then the pseudometric $D^{\mathcal{F}}$ is finite.*

Compared to other metrics and measurements on spaces of Lagrangians, the key novelty here is that properties (i), (ii), (iii), and (iv) are valid for the same metric.

Remark 3.5. (a) Point (i) of Theorem 3.4 shows that all the fragmentation pseudometrics $D^{\mathcal{F}}$ are dominated by the spectral metric. In previous results involving metrics on spaces of Lagrangians, such as those based on the shadows of cobordisms in [10], the best one could do was to establish upper bounds on the metrics that are generally much harder to estimate, such as the Hofer distance. Further consequences of this point will be discussed in Section 3.4.1.

(b) Point (ii) of Theorem 3.4 can be read as a typical non-squeezing-type result: embeddings of large symplectic balls, as in the definition of δ, are obstructed by $D^{\mathcal{F}}$. Conversely, this point implies that if $D^{\mathcal{F}}(L, L') = 0$, then $L \subset L' \cup \bigcup_{F \in \mathcal{F}} F$. As a result, suppose that we fix a second family $\mathcal{F}' \subset \mathcal{X}$, obtained through a small Hamiltonian perturbation of the elements of $\mathcal{F}$. One can then consider $D^{\mathcal{F},\mathcal{F}'} = \max\{D^{\mathcal{F}}, D^{\mathcal{F}'}\}$ as in (2.42). This pseudometric is non-degenerate on $\overline{\mathcal{X}}$ in the sense that $D^{\mathcal{F},\mathcal{F}'}(L, L') = 0$ if and only if $\overline{L} = \overline{L}'$ (in other words, the two underlying Lagrangians involved coincide; obviously, the markings may differ). This type of argument first appeared in [10]. Various forms of the inequality in point (ii) appeared earlier in the literature, in particular in cases such as when $\mathcal{F} = \{0\}$, where the metric involved is the Hofer metric (see, e.g., [4]). However, it is useful to note that the pseudometrics $D^{\mathcal{F}}$ are, in general, smaller compared to the metrics in these earlier references. Note also that, even for $\mathcal{F} = \{0\}$, the inequalities obtained by combining (i) and (ii) appear to be new.

(c) We emphasize that in point (iii) of Theorem 3.4 the two Lagrangians L and L' are allowed to be very different. For instance, they can be in different smooth isotopy classes or even have different homeomorphism types, and still $D^{\mathcal{F}}(L, L')$ can be finite (this point is reinforced by the last part of the theorem). Therefore, this result shows a form of rigidity of Lagrangian intersections for perturbations that are small in this metric $D^{\mathcal{F}}$, but that can be very large (even infinite) in other metrics. The result extends earlier persistence-type statements in Morse and Floer theory (one of the earliest examples appearing in [23]), most of them expressed in terms of the Hofer distance, which is much larger than $D^{\mathcal{F}}$. Again, there is considerable interest in working with the algebraic metrics introduced here because for other metrics, such as the shadow metrics based on Lagrangian cobordism, the finiteness result in point (iv) is not known to hold.

(d) The fact that we have $j_{\mathcal{X}}^* \mathcal{Y}(N) \in \mathrm{Obj}(\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty)$ for each marked Lagrangian N that intersects transversely the elements of $\mathcal{X}$, as in Theorem 3.1, implies that we can define a pseudometric on the space of all marked Lagrangians in X by

$$\Delta^{\mathcal{F}}(N, N') = \limsup_{\epsilon \to 0} \widehat{d^{\mathcal{F}}}\big(j_{\mathcal{X}}^* \mathcal{Y}(N_\epsilon), j_{\mathcal{X}}^* \mathcal{Y}(N'_\epsilon)\big),$$

where N_ϵ and N'_ϵ are ϵ-small (in the Hofer metric) Hamiltonian perturbations of N and N', respectively, both transverse to the elements of $\mathcal{X}$. This pseudometric is in general degenerate, as it does not "see" differences between N and N' away from the elements of $\mathcal{X}$.

(e) The constants providing the various bounds in Theorem 3.4 are very rough and can be improved in some cases, but we will not pursue this question here.

We will prove a consequence of Theorem 3.4, deduced by studying how the pseudometrics $D^{\mathcal{F}}$ change when the underlying set of marked Lagrangians $\mathcal{X}$ changes.

To state this consequence we need a global finiteness-type assumption on our Liouville manifold (X, ω). To formulate it, we denote by $\mathcal{L}ag(X)$ the set of exact, compact, graded, embedded Lagrangians in X, and by $\mathcal{L}ag(X)'$ the marked, exact Lagrangians in X (these are the elements of $\mathcal{L}ag(X)$ but with fixed primitives and grading choices). As before, the Fukaya category $D\mathcal{F}uk(X)$ is the derived category of the A_∞-category with objects the elements in $\mathcal{L}ag(X)'$. The category $D\mathcal{F}uk(X)$ is constructed as in [53]. In particular, the perturbation data depends only on the elements in $\mathcal{L}ag(X)$, and not on the choices of primitives and grading.

Definition 3.6. Let (X, ω) be a Liouville manifold. The *Fukaya rank* of (X, ω), rank $\mathcal{F}uk(X, \omega)$, is the minimal cardinality of a family of Lagrangians $\overline{\mathcal{F}} \subset \mathcal{L}ag(X)$ such that the corresponding family of marked Lagrangians $\mathcal{F} \subset \mathcal{L}ag(X)'$, obtained from $\overline{\mathcal{F}}$ by adding all possible translates of the objects in $\mathcal{F}$ (in terms of grading), generates $D\mathcal{F}uk(X)$. (Note that the primitives have no effect here, since we are talking about generators in a non-filtered setting.)

We emphasize here that "generating" has the meaning of triangulated generation, as everywhere else in this memoir. The rank $\mathcal{D}$ can be defined similarly for any triangulated category $\mathcal{D}$. The terminology is justified by the fact that this quantity is an upper bound for the rank of the Grothendieck group $K(\mathcal{D})$.

Corollary 3.7. *Let (X,ω) be a Liouville manifold. Assume that* $\operatorname{rank} \mathcal{F}uk(X,\omega) < \infty$. *Fix a family of generators $\mathcal{F} \subset \mathcal{L}ag(X)'$, invariant under shifts and translations, and such that the corresponding family $\overline{\mathcal{F}} \subset \mathcal{L}ag(X)$ obtained by forgetting the markings is finite and is in general position (in the sense defined at the beginning of Section 3.1). Then the set $\mathcal{L}ag(X)$ carries a finite pseudometric $\mathcal{D}^{\mathcal{F}}$ such that:*

(i) *(Spectrality) For any $L, L' \in \mathcal{L}ag(X)$, we have*

$$\mathcal{D}^{\mathcal{F}}(L,L') \leq 4\,\sigma(L,L').$$

(ii) *(Non-degeneracy) If $L, L' \in \mathcal{L}ag(X)$, then*

$$\frac{\delta(L; L' \cup \bigcup_{F\in\mathcal{F}} F)}{8} \leq \mathcal{D}^{\mathcal{F}}(L,L').$$

(iii) *(Persistence of intersections) Assume that $L, L', N \in \mathcal{L}ag(X)$ are in general position and $L' \notin \mathcal{F}$. If*

$$\mathcal{D}^{\mathcal{F}}(L,L') < \frac{1}{16}\,\delta^{\cap}\big(N, L'; \textstyle\bigcup_{F\in\mathcal{F}} F\big),$$

then

$$\#(L \cap N) \geq \#(L' \cap N).$$

In particular, if $\mathcal{F}'$ is another family obtained from $\mathcal{F}$ by generic Hamiltonian perturbations of the elements of $\overline{\mathcal{F}}$, then

$$\mathcal{D}^{\mathcal{F},\mathcal{F}'} = \max\{\mathcal{D}^{\mathcal{F}}, \mathcal{D}^{\mathcal{F}'}\}$$

is a finite and non-degenerate metric on $\mathcal{L}ag(X)$ that satisfies the properties (i), (ii), and (iii) above.

Thus, under the hypothesis $\operatorname{rank} \mathcal{F}uk(X,\omega) < \infty$, *the set $\mathcal{L}ag(M)$ has a metric space structure with respect to a metric satisfying the properties (i), (ii), and (iii).*

3.2 Filtrations in Floer homology and Fukaya categories

3.2.1 Filtered A_∞-categories and their associated TPCs

A filtered A_∞-category $\mathcal{A}$ is an A_∞-category over a given base field $\mathbf{k}$, such that the spaces of morphisms $\hom_{\mathcal{A}}(X,Y)$ between every two objects X, Y are filtered

(with increasing filtrations) and *all* the composition maps μ_d, $d \geq 1$, respect the filtrations. We endow $\hom_{\mathcal{A}}(X, Y)$ with the differential μ_1 and view it as a filtered chain complex. We denote by $\hom^s_{\mathcal{A}}(X, Y)$, $s \in \mathbb{R}$, the level-s filtration subcomplex of $\hom_{\mathcal{A}}(X, Y)$. We refer the reader to [10] for more details on filtered A_∞-categories. Note, however, that in [10] the theory is developed for the more general case of *weakly* filtered A_∞-categories (the "genuinely" filtered case is obtained from the weakly filtered one by assuming the so-called "discrepancies" of $\mathcal{A}$, defined in [10], to vanish). Of course, the A_∞ considerations here are very similar to those for dg-categories in Section 2.5.1.

For simplicity, we make three further assumptions on our filtered A_∞-categories. The first is that $\mathcal{A}$ is strictly unital, with the units lying in persistence level 0. The second one is that for every two objects $X, Y \in \mathrm{Obj}(\mathcal{A})$, the space $\hom_{\mathcal{A}}(X, Y)$ is finite-dimensional over $\mathbf{k}$. The third assumption is that $\mathcal{A}$ is complete with respect to persistence shifts, in the sense that we have a shift "functor" consisting of a family of A_∞-functors $\Sigma = \{\Sigma^r : \mathcal{A} \to \mathcal{A}, r \in \mathbb{R}\}$ whose members satisfy the following conditions:

(1) Σ^r is strictly unital and the higher components $(\Sigma^r)_d$, $d \geq 2$, of Σ^r all vanish.

(2) $\Sigma^0 = \mathbb{1}$, $\Sigma^s \circ \Sigma^t = \Sigma^{s+t}$.

(3) We are given prescribed identifications $\hom^s_{\mathcal{A}}(\Sigma^r X, Y) \cong \hom^{s+r}_{\mathcal{A}}(X, Y)$ that are compatible with the inclusions $\hom^\alpha_{\mathcal{A}}(X, Y) \subset \hom^\beta_{\mathcal{A}}(X, Y)$ for $\alpha \leq \beta$. These identifications are considered as part of the structure of the shift functor Σ.

The assumption that $\mathcal{A}$ is complete with respect to shifts is merely a matter of convenience, in the sense that it is not essential to impose this condition in advance. Indeed, any filtered A_∞-category (satisfying all the above assumptions except the one on completeness with respect to shifts) can be completed with respect to shifts by adding suitable objects that will play the role of the shifted $\Sigma^r X$ objects and then defining the functors Σ^r accordingly. See again Section 2.5.1 for the similar case of filtered dg-categories.

We will generally use homological conventions in the context of A_∞-categories; however, for compatibility with the literature we will generally use cohomological grading. Whenever this is the case, we will denote the cohomological degrees by superscripts (e.g., H^0 will stand for the homology in cohomological degree 0, the units will be assumed to be in cohomological degree 0, and so on).

Given a filtered A_∞-category $\mathcal{A}$, one can form the category $\mathrm{Tw}(\mathcal{A})$ of twisted complexes over $\mathcal{A}$, which is itself a filtered A_∞-category (satisfying all the additional assumptions mentioned earlier). This can be done by following the construction in [53, Chapter I, Section (3l)] and extending the filtrations from $\mathcal{A}$ to $\mathrm{Tw}(\mathcal{A})$ in the obvious way. The construction of the filtered $\mathrm{Tw}(\mathcal{A})$ in the case of dg-categories has

been worked out in detail in Section 2.5.1, and the A_∞ case is very similar. There is a bit of abuse of notation in writing $\mathrm{Tw}(\mathcal{A})$, since the latter category carries additional structures (namely filtrations and a shift functor) beyond those of the unfiltered category of twisted complexes, which is denoted in the literature by the same notation $\mathrm{Tw}(\mathcal{A})$.

The filtered A_∞-category $\mathcal{A}$ embeds into $\mathrm{Tw}(\mathcal{A})$ in an obvious way, the embedding being a filtered A_∞-functor which is full and faithful (on the chain level). Moreover, $\mathrm{Tw}(\mathcal{A})$ is pre-triangulated in the filtered sense (which in particular means that it is closed under formation of filtered mapping cones). It follows that the homological category $H^0(\mathrm{Tw}(\mathcal{A}))$ is a TPC that contains the homological persistence category $H^0(\mathcal{A})$ of $\mathcal{A}$.

Another TPC associated with the A_∞-category $\mathcal{A}$ is provided by the category $\mathrm{Fmod}(\mathcal{A})$ of filtered A_∞-modules over $\mathcal{A}$. Weakly filtered modules are defined in [10] and the filtered definitions correspond to all discrepancies being 0. We will only consider strictly unital modules here. There is a natural shift functor on this category, $\Sigma : (\mathbb{R}, +) \to \mathrm{End}(\mathrm{Fmod}(\mathcal{A}))$. Given $r \in \mathbb{R}$ and a module $\mathcal{M} \in \mathrm{Fmod}(\mathcal{A})$, we define the filtered module $\Sigma^r \mathcal{M}$ by $(\Sigma^r \mathcal{M})^{\leq \alpha}(N) = \mathcal{M}^{\alpha - r}(N)$, endowed with the same μ_d-operations as $\mathcal{M}$.

Remark 3.8. In the cases of interest to us, namely the Fukaya category, this shift functor on $\mathrm{Fmod}(\mathcal{A})$ is compatible with a shift functor on $\mathcal{A}$.

The category $\mathrm{Fmod}(\mathcal{A})$ is in fact a filtered dg-category in the sense of Section 2.5.1 and it is pre-triangulated. Thus $H^0(\mathrm{Fmod}(\mathcal{A}))$ is a TPC, by Corollary 2.102. Of more interest to us is a subcategory of $\mathrm{Fmod}(\mathcal{A})$. First notice that, because $\mathcal{A}$ is filtered, the Yoneda functor $\mathcal{Y} : \mathcal{A} \to \mathrm{Fmod}(\mathcal{A})$ is filtered too. Moreover, our assumption of strict unitality of $\mathcal{A}$ implies that $\mathcal{Y}$ is homologically full and faithful. Furthermore, there exists a canonical map

$$\lambda : \mathcal{M}(X) \to \hom_{\mathrm{Fmod}}(\mathcal{Y}(X), \mathcal{M})$$

for all $X \in \mathrm{Obj}(\mathcal{A})$ and $\mathcal{M} \in \mathrm{Obj}(\mathrm{Fmod}(\mathcal{A}))$, as defined in [53, Chapter I, Section (11)]. Standard arguments show that λ is a *filtered quasi-isomorphism*, in the sense that it is filtered and induces an isomorphism between the persistence homologies of its domain and target filtered chain complexes.

We consider now the pre-triangulated closure $\mathcal{A}^\#$ of the Yoneda modules and their shifts: this is a full subcategory of $\mathrm{Fmod}(\mathcal{A})$ that has as objects all the iterated cones, over filtration-preserving morphisms, of shifts of Yoneda modules (thus of modules of the form $\Sigma^r \mathcal{Y}(X)$).

Finally, we denote by $\mathcal{A}^\nabla$ the smallest full subcategory of $\mathrm{Fmod}(\mathcal{A})$ that contains $\mathcal{A}^\#$ and all the modules (and all their shifts and translates) that are r-quasi-isomorphic to objects in $\mathcal{A}^\#$ for some $r \in [0, \infty)$. Here, a module $\mathcal{M}$ is called r-quasi-isomorphic to $\mathcal{N}$ if, in $H^0(\mathrm{Fmod}(\mathcal{A}))$, there is an r-isomorphism $\mathcal{M} \to \mathcal{N}$.

It is easy to see that $\mathcal{A}^\nabla$ remains pre-triangulated, carries the shift functor induced from $\mathrm{Fmod}(\mathcal{A})$, and thus $H^0(\mathcal{A}^\nabla)$ is a TPC.

3.2.2 Persistence Floer homology

3.2.2.1 Filtered Floer complexes. Given a pair of marked Lagrangians L_0, L_1 as above and a choice of Floer data $\mathcal{D}_{L_0,L_1} = (H_{L_0,L_1}, J_{L_0,L_1})$, which consists of a (possibly time-dependent) Hamiltonian function and a choice of a compatible (time-dependent) almost complex structure, we can form the Floer complex $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$. This is a $\mathbb{Z}_2$-graded chain complex (recall our Lagrangians are assumed to be oriented). It is generated by the Hamiltonian chords $x : [0,1] \to X$ of H_{L_0,L_1} with end points on the two Lagrangians, namely $x(0) \in L_0$, $x(1) \in L_1$. For simplicity, we work here with coefficients in $\mathbb{Z}_2$.

Moreover, $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is a filtered chain complex, where the filtration function is given by the action functional. More precisely, if $x \in \mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is a generator (i.e., a Hamiltonian chord), its action is defined by

$$\mathcal{A}(x) := \int_0^1 H_{L_0,L_1}(t, x(t))\, dt - \int_0^1 \lambda(\dot{x}(t))\, dt + h_{L_1}(x(1)) - h_{L_0}(x(0)).$$

Remark 3.9. In case L_0 and L_1 intersect transversely and $H_{L_0,L_1} \equiv 0$, the Hamiltonian chords x that generate $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ are just the intersection points $L_0 \cap L_1$ and the action reduces to

$$\mathcal{A}(x) = h_{L_1}(x) - h_{L_0}(x) \quad \forall\, x \in L_0 \cap L_1. \tag{3.3}$$

Back to the general case, the homology of the filtered chain complex $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ gives rise to the persistence Floer homology $\mathrm{HF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$, which has the structure of a $\mathbb{Z}_2$-graded persistence module (over the field $\mathbb{Z}_2$). As a vector space, $\mathrm{HF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is independent of the auxiliary Floer data $\mathcal{D}_{L_0,L_1}$; however, as a persistence module, it does depend on that choice. More precisely, the persistence module structure of $\mathrm{HF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is independent of the choice of the almost complex structure J_{L_0,L_1} from $\mathcal{D}_{L_0,L_1}$; however, it depends strongly on the choice of the Hamiltonian H_{L_0,L_1}.

3.2.2.2 Grading. While $\mathbb{Z}_2$-grading is enough for our applications, one can obtain a $\mathbb{Z}$-graded theory if one makes additional assumptions on X and on the admissible class of Lagrangian submanifolds. The simplest such conditions are the following. Firstly, we assume that $2c_1(X) = 0$, where $c_1(X)$ stands for the first Chern class of the tangent bundle of X, viewed as a complex vector bundle by endowing X with any ω-compatible almost complex structure J. We now fix a nowhere vanishing quadratic complex n-form (where $n = \dim_{\mathbb{C}} X$), namely a nowhere vanishing section Θ of the bundle $\Omega^n(X, J)^{\otimes 2}$. The choice of Θ gives rise to a global phase map

$\det^2_\Theta : \mathcal{L}(T(X)) \to S^1$ defined on the Lagrangian Grassmannian bundle $\mathcal{L}(T(X))$ of X (see [52], [53, Chapter II, Section (11j)]). Given a Lagrangian $\bar{L} \subset X$, denote by $s_{\bar{L}} : \bar{L} \to \mathcal{L}(T(X))|_{\bar{L}}$ its Gauss map. A Lagrangian $\bar{L}$ is said to admit a grading if $\det^2_\Theta \circ s_{\bar{L}} : \bar{L} \to S^1$ can be lifted to a function $\theta_{\bar{L}} : \bar{L} \to \mathbb{R}$, and a choice of such a lift is called a grading on $\bar{L}$. In this case, by adding integral constants to $\theta_{\bar{L}}$ one obtains all possible gradings of $\bar{L}$.

Gradability of Lagrangians can be rephrased in cohomological terms. The map $\det^2_\Theta \circ s_{\bar{L}} : \bar{L} \to S^1$ gives rise to a cohomology class $\mu_{\bar{L}} \in H^1(\bar{L})$, which we call the Maslov class of $\bar{L}$. (There is a slight abuse of notation here since $\mu_{\bar{L}}$ actually depends on the homotopy class of Θ.) A Lagrangian $\bar{L}$ admits a grading if and only if $\mu_{\bar{L}} = 0$.

The relation between $\mu_{\bar{L}}$ and the more familiar Maslov index homomorphism $\mu_{X,\bar{L}} : H_2(X, \bar{L}) \to \mathbb{Z}$ is that $\mu_{X,\bar{L}}(A) = \langle \mu_{\bar{L}}, \partial_* A \rangle$ for every $A \in H_2(X, \bar{L})$, where $\partial_* : H_2(X, \bar{L}) \to H_1(\bar{L})$ is the connecting homomorphism. Note also that if the map $H_1(\bar{L}) \to H_1(X)$, induced by the inclusion $\bar{L} \subset X$, is trivial, then $\mu_{\bar{L}}$ is determined by $\mu_{X,\bar{L}}$ (hence in that case $\mu_{\bar{L}}$ is independent of the choice of Θ). This is because $\mu_{X,\bar{L}}(j(B)) = 2\langle c_1(X), B \rangle = 0$ for every $B \in H_2(B)$, where the map $j : H_2(X) \to H_2(X, \bar{L})$ is induced by the inclusion. Therefore $\mu_{X,\bar{L}}$ descends to $H_2(X, \bar{L})/j(H_2(X)) \cong H_1(\bar{L})$.

In the rest of the memoir, we optionally allow for a $\mathbb{Z}$-graded theory. Whenever this is desired, we will make the preceding assumptions on X, fix the auxiliary structure Θ, and consider only marked Lagrangians L that admit a grading (or, equivalently, $\mu_{\bar{L}} = 0$). Moreover, we extend the notion of a marked Lagrangian L to include also a choice of a grading denoted by θ_L, namely $L = (\bar{L}, h_L, \theta_L)$. However, below we will mostly suppress the choice θ_L from the notation, since it will not often be used explicitly.

Given a pair of marked Lagrangians L_0, L_1, their gradings induce an absolute $\mathbb{Z}$-grading on $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$, and therefore also on $\mathrm{HF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$. The effect of translating the grading functions on the Lagrangians is the following. Write $L[k] = (\bar{L}, h_L, \theta_L - k)$. Then, using cohomological and homological grading respectively, we have

$$\mathrm{CF}^i(L_0[k], L_1[l]; \mathcal{D}) \cong \mathrm{CF}^{i+k-l}(L_0, L_1; \mathcal{D}),$$
$$\mathrm{CF}_j(L_0[k], L_1[l]; \mathcal{D}) \cong \mathrm{CF}_{j+l-k}(L_0, L_1; \mathcal{D}).$$

3.2.3 Weakly filtered Fukaya categories

The above construction can be enhanced to an A_∞-category called the Fukaya category.

Fix a collection of marked Lagrangians $\mathcal{X}$ in X. We assume that $\mathcal{X}$ is closed under grading translations and shifts of the primitives, namely if $L = (\bar{L}, h_L, \theta_L)$ is in $\mathcal{X}$, then for every $r \in \mathbb{R}$ and $k \in \mathbb{Z}$, the marked Lagrangian $\Sigma^r L[k] := (\bar{L}, h_L + r,$

$\theta_L - k$) is also in $\mathcal{X}$. (Of course, in case $\mathcal{X}$ is not closed under shifts and translations we can easily fix this by adding to $\mathcal{X}$ all the shifts and translations of its objects.)

The Fukaya category $\mathcal{Fuk}(\mathcal{X})$ associated to $\mathcal{X}$ is an A_∞-category whose objects are the elements of $\mathcal{X}$ and the complex of morphisms between a pair of objects from $\mathcal{X}$ is the Floer complex of that pair. In order to set up this A_∞-category one has to choose for every pair of objects (L_0, L_1) from $\mathcal{X}$ a regular Floer datum $\mathcal{D}_{L_0,L_1}$ and then extend this choice to a consistent choice of regular perturbation data $\mathcal{P}_{\mathcal{X}}$, which is defined for every tuple of Lagrangians $(L_0, \ldots, L_d)$, $d \geq 1$, from the collection $\mathcal{X}$. (It is important that both the Floer data as well as the perturbation data associated to a tuple depend only on the underlying Lagrangians in that tuple, and not on the choice of primitives or gradings on the Lagrangians in the tuple.)

Once these choices are set, one defines

$$\hom_{\mathcal{Fuk}(\mathcal{X})}(L_0, L_1) := \mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1}),$$

endowed with the Floer differential μ_1. The higher-order operations μ_d for $d \geq 2$ are multilinear maps

$$\mu_d : \mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1}) \otimes \cdots \otimes \mathrm{CF}(L_{d-1}, L_d; \mathcal{D}_{L_{d-1},L_d}) \to \mathrm{CF}(L_0, L_d; \mathcal{D}_{L_0,L_d})$$

of cohomological degree $2 - d$, which are defined for every tuple of Lagrangians $L_0, \ldots, L_d$ from $\mathcal{X}$. They satisfy the A_∞-identities. The definition of μ_d goes by counting Floer $(d + 1)$-polygons in X with boundary conditions on the L_i's. These polygons satisfy a perturbed Cauchy–Riemann equation with perturbations prescribed by $\mathcal{P}_{\mathcal{X}}$. Note that the Fukaya category described above depends on the choice $\mathcal{P}_{\mathcal{X}}$ of the perturbation data, hence should in fact be denoted by $\mathcal{Fuk}(\mathcal{X}; \mathcal{P}_{\mathcal{X}})$. However, it is well known that different choices of perturbation data lead to quasi-equivalent categories [53, Chapter II, Section 10]

Taking filtrations into account, as already mentioned in Section 3.2.2.1, the hom's of this category are filtered chain complexes. However, due to the perturbation data involved in defining the higher-order operations, the μ_d-operations for $d \geq 2$ do not preserve the action filtrations, but only do so up to an error (that depends on d). Consequently, the resulting A_∞-category is not filtered but only *weakly filtered*. Enhancing such a structure to a TPC, e.g., along the lines of the construction outlined in Section 3.2.1, seems like a non-trivial technical problem.

3.3 Genuinely filtered Fukaya categories

Here we will outline a construction that gives rise to a genuinely filtered Fukaya A_∞-category. This, however, will require very restrictive assumptions on the collection of objects $\mathcal{X}$, and some adjustments in the definition of the operations μ_d for certain tuples of Lagrangians.

Recall from the beginning of the section that we denote by $\overline{\mathcal{X}} = \{\overline{L} \mid L \in \mathcal{X}\}$ the collection of underlying Lagrangian submanifolds corresponding to the marked Lagrangians in $\mathcal{X}$. Recall also the assumptions on $\overline{\mathcal{X}}$: $\overline{\mathcal{X}}$ is finite; every two distinct Lagrangians $L', L'' \in \overline{\mathcal{X}}$ intersect transversely; for every three distinct Lagrangians $\overline{L}_0, \overline{L}_1, \overline{L}_2 \in \overline{\mathcal{X}}$ we have $\overline{L}_0 \cap \overline{L}_1 \cap \overline{L}_2 = \emptyset$. We also continue to assume, as before, that $\mathcal{X}$ is closed under shifts and translation in grading.

A more general approach, yielding genuinely filtered Fukaya categories, has been developed by Ambrosioni [1] after the first version of the current work had appeared. This approach does not impose any restrictions on the collection $\mathcal{X}$, besides the assumption that all Lagrangians in $\mathcal{X}$ are weakly exact (and possibly graded, if one wants a graded theory). In particular, no finiteness condition on $\overline{\mathcal{X}}$ is needed in that work, and no transversality assumption is made for distinct elements of this collection. On the other hand, the invariance properties of the filtered Fukaya categories from [1] are coarser than the ones provided by our approach (compare, e.g., Theorems 3.1 and 3.12 to [1, Theorem B]).

The construction outlined below is based on methods already well established in the literature, and we will therefore only provide a sketch of the construction, omitting quite a few technical details but emphasizing some points that are important for the control of filtration aspects.

3.3.1 Floer chain complexes redefined

We begin by redefining the Floer chain complexes in a way that will enable us to obtain a genuinely filtered Fukaya category.

Let $L_0, L_1 \in \mathcal{X}$. Assume first that $\overline{L}_0 \neq \overline{L}_1$ (hence they intersect transversely). In this case, we fix a Floer datum $\mathcal{D}_{L_0,L_1}$ of the type $(0, J_{L_0,L_1})$, i.e., its Hamiltonian term will be identically 0. (Once again, the choice of J_{L_0,L_1} is made such that it depends only on the underlying Lagrangians $\overline{L}_0$, $\overline{L}_1$.) We then define $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ to be the standard Floer complex associated to the pair $(\overline{L}_0, \overline{L}_1)$ using the Floer data $\mathcal{D}_{L_0,L_1}$ chosen above. The grading is defined using the grading on the two marked Lagrangians L_0, L_1. The filtration on $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is defined by using the action as a filtration function. Specifically, if $x \in \overline{L}_0 \cap \overline{L}_1$ is a generator of $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$, its action $\mathcal{A}(x)$ is defined by (3.3).

Assume now that $\overline{L}_0 = \overline{L}_1$ and denote by $\overline{L}$ this common Lagrangian. In this case, the Floer datum will be replaced by a choice of a Morse datum (which we continue to denote by $\mathcal{D}_{L_0,L_1}$), namely a pair $(f_{\overline{L}}, (\cdot,\cdot)_{\overline{L}})$ of a Morse function $f_{\overline{L}} : \overline{L} \to \mathbb{R}$ and a Riemannian metric $(\cdot,\cdot)_{\overline{L}}$ on the common underlying Lagrangian $\overline{L}$. We will further assume that all the Morse functions $f_{\overline{L}}$ have a unique (local) maximum (i.e., a unique critical point of index $n = \dim_{\mathbb{C}} X$). The purpose of this assumption is to ensure that the units in our Fukaya category will be strict (rather than only homology units).

The Floer complex $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is defined to be the Morse complex $\mathrm{CM}(\bar{L})$ of $\bar{L}$, associated to the Morse data $\mathcal{D}_{L_0,L_1} = (f_{\bar{L}}, (\cdot,\cdot)_{\bar{L}})$. We filter this chain complex in the following way. We set the filtration level *for all* generators $x \in \mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ of this chain complex (which are critical points of $f_{\bar{L}}$) to be the constant $c \in \mathbb{R}$, where $c \equiv h_{L_1} - h_{L_0}$ is the difference of the primitive functions of the two markings of the Lagrangian $\bar{L}$. To keep notation uniform, we continue to denote the filtration level of x by $\mathcal{A}(x)$, in the same way as we have done for the action.

3.3.2 Clusters of punctured disks

To define the μ_d-operations we will use a hybrid model that combines Floer polygons with gradient Morse trees. The maps defining the μ_d-operations will be called clusters of Floer polygons. This approach is analogous to the cluster Floer homology theory initiated by Cornea–Lalonde [21], who also introduced the name "clusters" in this context. Further modifications and foundational work on the subject has been done in [15, 16]. The main difference between these works and what we will be doing below is the following. The cluster homology theory [21] deals with a single Lagrangian in the presence of pseudoholomorphic disks. The "clusters" in that work consist of Morse flow lines attached to pseudoholomorphic disks. In contrast, here we deal with Floer theory of many Lagrangians together (setting up a Fukaya category), but in the absence of pseudoholomorphic disks. Note that [1] too uses clusters of Floer polygons in much the same way as below, and contains a detailed account of the subject.

In order to describe clusters of Floer polygons, we first need to set up their domains, which we call clusters of punctured disks.

We begin with the notion of a *k-punctured disk*. By this we mean a Riemann surface S_k which is obtained from the closed 2-disk $D \subset \mathbb{C}$ by removing $k \geq 1$ distinct boundary points $z_1, \ldots, z_k \in \partial D$, ordered in the *clockwise* direction, together with the following additional data. The points z_i will be called punctures. Each puncture z_i is declared to be either an *entry* puncture or an *exit* puncture. We allow S_k to have at most one exit puncture. We will typically denote the entry punctures by a $+$ superscript (e.g., z^+) and the exit puncture by a $-$ superscript (e.g., z^-). See Figure 3.1. Note that the boundary ∂S_k consists of k arcs, which we typically denote by $C_1, \ldots, C_k$, where the arc C_j goes from z_j to z_{j+1} for $1 \leq j \leq k-1$, in the clockwise direction, and C_k goes from z_k to z_1.

For each punctured disk S_k we fix a choice of strip-like ends along each of its punctures, as in [53, Chapter II, Section 9]. These choices should be compatible with splitting and gluing, as will be described later on.

Next we need to consider certain types of trees, which we call admissible. An admissible tree T is a (connected) tree with a finite number of edges and with the

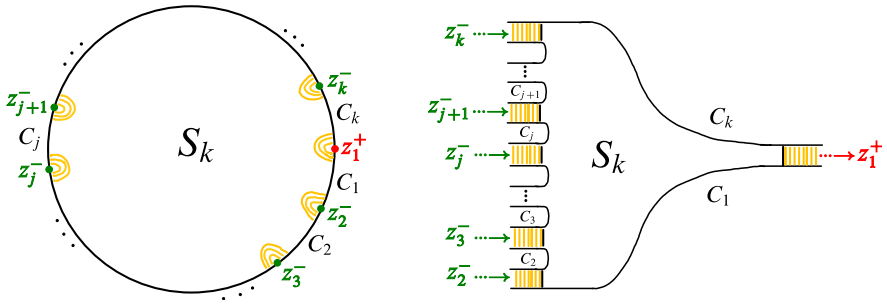


Figure 3.1. On the left: a k-punctured disk S_k with $k-1$ entries and one exit. The regions in yellow are the strip-like ends. On the right: a Riemann surface biholomorphic to a k-punctured disk, illustrating the strip-like ends modeled on $(-\infty, 0] \times [0, 1]$ for the entries and $[0, \infty) \times [0, 1]$ for the exit.

following properties and additional structures. In what follows, we will call all the end-vertices of T *leaves* (in particular, we will not distinguish between a possible root of the tree and the other end-vertices, and will just refer to all of them simply as leaves). We assume that all the leaves have valency 1 and all the other vertices of T (i.e., the internal ones) have valency 3. Moreover, the edges of T are oriented and these orientations satisfy the property that at every internal vertex (which by assumption has valency 3) there are precisely two incoming edges and one outgoing edge. The leaves of the tree T are divided into two types: E and A, where E stands for entry/exit leaves and A for attachment leaves. A leaf of type E will be called an *entry leaf* if the orientation on the edge connected to it goes from the leaf towards the rest of the tree. In the opposite case, i.e., when the orientation of that edge goes into the leaf, it will be called an *exit leaf*. The edges of T that are not connected to type-E leaves will be called *internal edges*. These consist of all edges that are not connected to any leaf, as well as those edges that are connected to leaves of type A. The other edges will be called external edges.

The edges of the trees are labeled by intervals in $\mathbb{R}$ as follows. The internal edges are labeled by intervals of the type $[0, R]$ (with possibly different values of $R > 0$ for different edges). The edges that connect a leaf of type E to an internal vertex are labeled either by $(-\infty, 0]$ or $[0, \infty)$, according to whether that leaf is an entry or an exit, respectively. If there is an edge connecting two leaves of type E (which happens if and only if the tree consists of exactly these two leaves and one edge connecting them), then this edge is labeled by the interval $(-\infty, \infty)$. Finally, we also fix an isotopy class of planar embeddings for the tree T. Note that, as a result, this fixes a cyclic clockwise order on the three edges connected to any given internal vertex. It also induces a cyclic clockwise order on the leaves of the entire tree. We illustrate a typical example of an admissible tree in Figure 3.2 below.

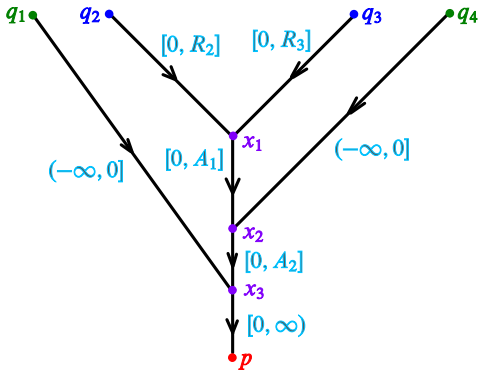


Figure 3.2. An admissible tree. The interval labelings are in dark gray. The leaves q_1, q_4 (in green) are type-E entry leaves, and p is a type-E exit leaf. The leaves q_2, q_3 (in blue) are of type A. The other vertices, x_1, x_2, x_3 (in purple), are interior vertices. The overall (cyclic) clockwise order on the leaves of this tree is q_1, q_2, q_3, q_4, p.

Having defined admissible trees, we now fix once and for all on each such tree T a collection of *orientation-preserving* identifications $\sigma_e^T : e \to I_e$ between each edge e and the interval I_e labeling it. Of course, in case the interval I_e is of the type $[0, \infty)$, $(-\infty, 0]$, or $(-\infty, \infty)$ (which happens when e is connected to vertices of type E) then the vertices corresponding to $\pm\infty$ are only asymptotically identified with $\pm\infty$. Note that every underlying tree has (infinitely) many different interval labelings for its edges (internal edges can be labeled by $[0, R]$ for different values of $R > 0$), leading to different admissible trees T. There is an obvious parametrization of these different interval labelings (basically by choosing the parameter R on each internal edge). We require the identifications σ^T to depend continuously on these parameters.

We are now in a position to introduce clusters of punctured disks. These are built from a collection $\mathcal{S} = \{S_{k_1}^{(1)}, \ldots, S_{k_l}^{(l)}\}$, $l \geq 0$, of punctured disks and a collection of admissible trees $\mathcal{T} = \{T_1, \ldots, T_r\}$, $r \geq 0$, which are attached to the punctured disks in $\mathcal{S}$ at their leaves of type A. The attachment of the trees is done as follows. Let $T \in \mathcal{T}$ and denote its leaves of type A by $a_1, \ldots, a_{s_T} \in T$. For each $1 \leq i \leq s_T$, we identify the point $a_i \in T$ with a point lying on the boundary $\partial S_{k_j}^{(j)}$ of one of the punctured disks from $\mathcal{S}$. Here $j = j(T, i)$ depends on the tree T and the index i of the vertex a_i that is being attached. These attachments are subject to the following rule: each type-A leaf of a given tree $T \in \mathcal{T}$ is attached to one, and only one, punctured disk, and no two type-A leaves of the same tree $T \in \mathcal{T}$ are attached to the same punctured disk. There is no type-A leaf from the trees in $\mathcal{T}$ that is left unattached. We also require that, among the type-A leaves of all the trees in $\mathcal{T}$, there are no two leaves that are attached to the same point on the boundaries of the punctured disks.

We denote the space resulting from the above attachments by

$$\Sigma = \left(\bigcup_{q=1}^{l} S_{k_q}^{(q)} \right) \cup_A \left(\bigcup_{p=1}^{r} T_p \right), \tag{3.4}$$

where $\cup_A$ stands for the attachments described earlier. We will denote the part of Σ coming from the punctured disks (i.e., the leftward union in the right-hand side of (3.4)) by Σ_S, and the part coming from the trees (the rightward union in the right-hand side of (3.4)) by Σ_{tr}.

We now impose further restrictions on the previously described attachments. Consider the space obtained from Σ by collapsing each punctured disk $S_{k_q}^{(q)}$ from Σ_S to a (different) point:

$$\widetilde{\Sigma} := \Sigma \big/ \big(S_{k_q}^{(q)} \sim \mathrm{point}_q \ \forall q \big). \tag{3.5}$$

We require that the attachments of the tree described above be made in such a way that $\widetilde{\Sigma}$ is path-connected and, moreover, is a tree (hence, in particular, simply connected). We do not require this tree to be admissible.

Going back again to the space Σ, we note that it comes with a set of distinguished points: the punctures of the disks $S_{k_q}^{(q)}$ together with the leaves of the trees $T \in \mathcal{T}$ that are of type E. We call these points *external points* and denote by Σ_{ep} the set of all such points. The total number of external points of Σ will be called the *order* of Σ.

The external points Σ_{ep} are divided into two types: "entry points" and "exit points", regardless of a point being a puncture or a type-E leaf of some tree. We require that Σ_{ep} has precisely one exit point (which can be either an exit leaf or an exit puncture). We also require that Σ_{ep} has at least one entry point. Finally, in case there are only two such points (i.e., one entry and one exit), we require that either Σ is just a disk punctured at two points (with no trees attached) or that Σ has no punctured disks at all and it consists of just a tree with two vertices and one edge connecting them.

The last requirement on Σ is the following. Consider the tree $\widetilde{\Sigma}$ defined in (3.5). Note that, by construction, each edge of this tree is oriented (since the edges of all $T \in \mathcal{T}$ are oriented). Moreover, this tree has a distinguished vertex P_{exit}, namely the vertex that corresponds either to the punctured disk $S_k \subset \Sigma_S$ that contains the unique exit puncture, or to the unique exit leaf that belongs to one of the trees of Σ_{tr}. We require that the orientation on the tree $\widetilde{\Sigma}$ has the property that given any vertex $p \in \widetilde{\Sigma}$ there is a path from p to the distinguished vertex P_{exit} that is compatible with the orientation on $\widetilde{\Sigma}$. Figure 3.3 illustrates two examples of clusters of punctured disks (the labeling by Lagrangians $\overline{L}_i$ of the trees and the arcs in the ∂S_j's that appear in the picture should be ignored for the moment; these will be explained later on).

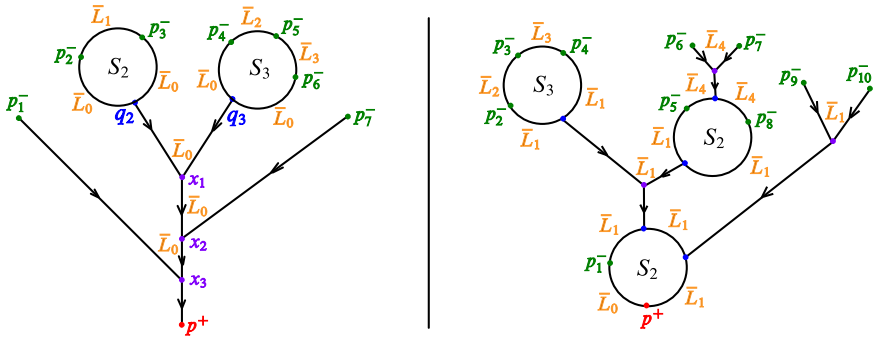


Figure 3.3. Two examples of clusters of punctured disks. The interval labeling of the edges of the trees is omitted here. The overall clockwise cyclic ordering of the external points on the left cluster is $(p_1^-, \dots, p_7^-, p^+)$ and on the right cluster $(p_1^-, \dots, p_{10}^-, p^+)$. The tuples describing the labeling by Lagrangians are $\bar{\mathcal{L}}_\Sigma = (\bar{L}_0, \bar{L}_0, \bar{L}_1, \bar{L}_0, \bar{L}_2, \bar{L}_3, \bar{L}_0, \bar{L}_0)$ for the left cluster and $\bar{\mathcal{L}}_\Sigma = (\bar{L}_0, \bar{L}_1, \bar{L}_2, \bar{L}_3, \bar{L}_1, \bar{L}_4, \bar{L}_4, \bar{L}_1, \bar{L}_1, \bar{L}_1)$ for the cluster on the right.

Let Σ, Σ' be two spaces obtained as above from two pairs of collections $\mathcal{S}$, $\mathcal{T}$ and $\mathcal{S}'$, $\mathcal{T}'$ of punctured disks and trees. We say that Σ and Σ' are equivalent if there is a homeomorphism $f : \Sigma \to \Sigma'$ with the following properties. The map f maps Σ_S biholomorphically to Σ'_S, and maps the trees Σ_{tr} to Σ'_{tr} by an isomorphism of trees (i.e., it maps vertices to vertices and edges to edges). Moreover, f intertwines all the other structures on $\Sigma_S, \Sigma_{\mathrm{tr}}$ with those on $\Sigma'_S, \Sigma'_{\mathrm{tr}}$. This means, in particular, that entry and exit punctures in Σ'_S correspond under f to the punctures of the same type in Σ_S; the same goes for the orientations on the edges of Σ_{tr}, Σ'_{tr}, the interval labeling, the identifications σ_e^T, and the classes of planar embeddings of the trees.

An equivalence class of spaces Σ as above (together with all the structures accompanying it) will be called a *cluster of punctured disks*. However, we will often use this name also for a specific representative Σ within a given equivalence class.

For a cluster of punctured disks, say represented by Σ, the orientation on the boundaries of the punctured disks in Σ_S and the classes of planar embeddings of the trees in Σ_{tr} induce a preferred clockwise cyclic order on the set of external points of Σ (recall that the external points consist of the entry and exit points, regardless of whether they are type-E leaves of the trees or punctures of the disks). Note that this ordering is preserved by the homeomorphisms defining the equivalence between different representatives Σ of the same class.

In what follows, it will be convenient to single out clusters of punctured disks of the following type. A cluster of punctured disks Σ is called *simple* if it consists of a single punctured disk without any trees attached.

We now turn to decorated clusters of punctured disks. Let Σ be a cluster of punctured disks. By a decoration of Σ by elements of $\bar{\mathcal{X}}$ we mean the following. We label

each arc in the boundaries of Σ_S, as well as each edge in the trees of Σ_{tr}, by an element of $\bar{\mathcal{X}}$. The labeling is subject to the following restrictions. In each tree from Σ_{tr} all the edges are labeled by the same $\bar{L} \in \bar{\mathcal{X}}$ (alternatively, one can think of each tree $T \subset \Sigma_{\mathrm{tr}}$ as being labeled by an element of $\bar{\mathcal{X}}$). The restriction on the labeling for the Σ_S-part of Σ is that in each punctured disk S_k from Σ_S there are no two consecutive arcs (i.e., two arcs with one puncture between them) that are labeled by the same element form $\bar{\mathcal{X}}$.

Once a cluster of punctured disks Σ is decorated by elements of $\bar{\mathcal{X}}$, we can form a tuple $\bar{\mathcal{L}}_\Sigma = (\bar{L}_0, \dots, \bar{L}_d)$ that encodes its decoration, where $d+1 = |\Sigma_{\mathrm{ep}}|$ is the order of Σ. The definition of $\bar{\mathcal{L}}_\Sigma$ goes as follows. Denote by $p^+, p_1^-, \dots, p_d^-$ the external points of Σ, ordered as explained earlier, where p^+ is the unique exit point and the p_j^-'s are all entry points. If p_j^- is a puncture of one of the disks $S_k \subset \Sigma_S$, we take $\bar{L}_j$ to be the Lagrangian that labels the arc on ∂S_k coming after the puncture (where "after" refers to the clockwise orientation on ∂S_k). If the entry p_j^- is a leaf of one of the trees $T \subset \Sigma_{\mathrm{tr}}$, then we take $\bar{L}_j$ to be the Lagrangian that labels that tree. We define $\bar{L}_0$ in the same way, according to whether p^+ is a puncture or a leaf. Figure 3.3 shows two examples of decorated clusters of punctured disks.

We will now reverse in some sense the decoration construction. Namely, we fix a tuple $\bar{\mathcal{L}} = (\bar{L}_0, \dots, \bar{L}_d)$ of Lagrangians from $\bar{\mathcal{X}}$, and consider the space $\mathrm{Clus}(\bar{\mathcal{L}})$ of all possible decorated clusters of punctured disks Σ with $\bar{\mathcal{L}}_\Sigma = \bar{\mathcal{L}}$. We call the elements of this space *$\bar{\mathcal{L}}$-decorated clusters of punctured disks*. As before, the elements of this space are equivalence classes of the spaces Σ, rather than the spaces Σ themselves. But it will often be convenient to work with an actual representative Σ of a given class.

Clearly every decorated cluster Σ belongs to a unique space $\mathrm{Clus}(\bar{\mathcal{L}})$, since the tuple $\bar{\mathcal{L}}_\Sigma$ is uniquely defined by Σ.

Remark 3.10. Let $\bar{\mathcal{L}} = (\bar{L}_0, \dots, \bar{L}_d)$ be a tuple of Lagrangians from $\bar{\mathcal{X}}$.

(a) If $\bar{\mathcal{L}}$ has the property that $\bar{L}_i \neq \bar{L}_j$ for every $i \neq j$, then every cluster of punctured disks Σ that admit an $\bar{\mathcal{L}}$-decoration must be simple.

(b) The converse statement to point (a) above is obviously not true whenever $d \geq 3$. Namely, one can decorate a simple cluster of punctured disks Σ by a tuple $\bar{\mathcal{L}}$ whose entries do have repetitions. However, in such a case we must have $\bar{L}_i \neq \bar{L}_{i+1}$ for every $0 \leq i \leq d$ (where the indexing is to be understood cyclically $\mathrm{mod}(d+1)$).

(c) A tuple $\bar{\mathcal{L}}$ with $\bar{L}_i \neq \bar{L}_{i+1}$ for every $0 \leq i \leq d$ can decorate also non-simple clusters of punctured disks. However, if a non-simple Σ is decorated by such an $\bar{\mathcal{L}}$ then none of the trees in Σ_{tr} can have external leaves (which means that each leaf in Σ_{tr} is attached to some punctured disk in Σ_S).

3.3.3 Splitting and degeneration

Given a tuple $\bar{\mathcal{L}} = (\bar{L}_0, \dots, \bar{L}_d)$ of Lagrangians from $\bar{\mathcal{X}}$, the space $\mathrm{Clus}(\bar{\mathcal{L}})$ of $\bar{\mathcal{L}}$-decorated clusters of punctured disks has the structure of a smooth manifold, analogous to the space of punctured disks from [53, Chapter II, Section 9]. This manifold admits a natural partial compactification into a manifold with corners. The top-dimensional strata of its boundary correspond to several types of degenerations of clusters of punctured disks, which we briefly describe below. Note that adding this boundary to $\mathrm{Clus}(\bar{\mathcal{L}})$ will still not make a full compactification of this space (hence the use of the words "partial compactification"); however, it will be enough for the purpose of establishing the A_∞-category identities. Below we will call those degenerations that lead to elements of this boundary *admissible degenerations* and their limiting objects *admissible degenerate clusters*.

Splitting within punctured disks. We begin by describing two variants of a degeneration that can occur to one punctured disk moving in a family. A family of punctured disks S^t_k (here k is the number of punctures and $t \in \mathbb{R}$ is parametrizing the family) can degenerate (or split) into two punctured disks $S'_{k'}$ and $S'_{k''}$, where $k' + k'' = k + 2$. The first punctured disk $S'_{k'}$ "inherits" $k' - 1$ of the entry punctures of S (placed in the same clockwise order as in S) and has one additional exit puncture z'_+. The other component, $S''_{k''}$, "inherits" all the other $k - (k' - 1) = k'' - 1$ punctures of S (again, in the same clockwise order) and has one additional entry puncture z''_-. Note that at the moment we do allow k' or k'' to take the values 1 or 2. This is in contrast to the more standard realizations of the Fukaya category, where each of the two disks in a splitting are required to have at least 3 punctures. However, later on, when viewing these disks as part of a degenerate cluster, more restrictions will be added in order to make such a degenerate configuration an admissible one.

Conversely, the two punctured disks $S'_{k'}$ and $S''_{k''}$ can be glued along the punctures z'_+, z''_- into a family of punctured disks S^t_k.

Depending on the context, in what follows we will sometimes view the preceding degeneration differently. Namely, regard the two punctures z'_+, z''_- as "removable" and view the degeneration of the family S^t_k as a splitting into two punctured disks $S'_{l'}$ and $S''_{l''}$ attached one to the other at a point (which is not a puncture) on their boundaries. Note that now we have $l' + l'' = k$, and similarly to the preceding case, we do allow l' or l'' to take the values 1 and 2.

Conversely, as before, the two punctured disks $S'_{l'}$ and $S''_{l''}$ can be glued into a family of punctured disks S^t_k.

Analytically, the two variants described above are the same; however, when taking decorations into account it is important to distinguish between them. More precisely, if the punctured disks S^t_k are decorated by the Lagrangians $(\bar{N}_0, \dots, \bar{N}_{k-1})$, the first variant of splitting corresponds to two punctured disks with decorations $(\bar{N}_r, \dots, \bar{N}_s)$

and $(\bar{N}_0, \dots, \bar{N}_r, \bar{N}_s, \dots, \bar{N}_{k-1})$, where $1 \le r < s \le k-1$, $\bar{N}_r \ne \bar{N}_s$, $s - r + 1 = k'$, and $k'' = k - (s - r) + 1$.

The second variant (i.e., where S^t_k degenerates into $S'_{l'}$ and $S''_{l''}$ attached at a common point, which is not a puncture, along their boundaries) corresponds to the case when the decoration $(\bar{N}_0, \dots, \bar{N}_{k-1})$ has $\bar{N}_r = \bar{N}_s$ for some non-consecutive indices, $r < s - 1$, and the splitting yields the decorations $(\bar{N}_r, \dots, \bar{N}_{s-1})$ and $(\bar{N}_0, \dots, \bar{N}_r, \bar{N}_{s+1}, \dots, \bar{N}_{k-1})$ on $S'_{l'}$ and $S''_{l''}$, respectively.

The two variants of splitting are illustrated in Figure 3.4.

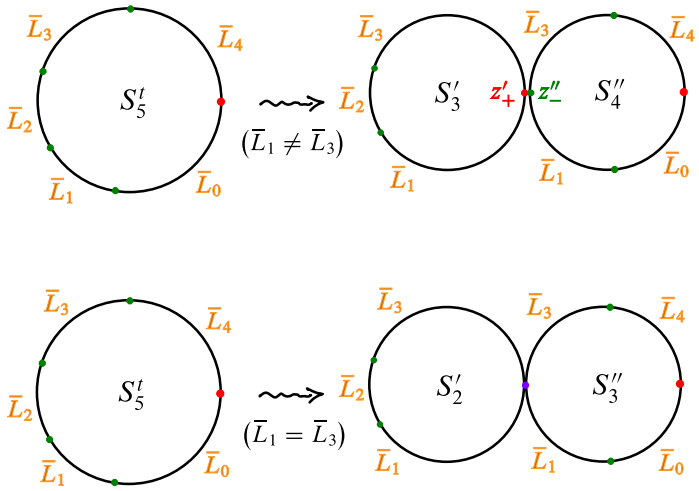


Figure 3.4. Two variants of splitting of decorated punctured disks. The points z'_+ (red) and z''_- (green) on the right-hand side of the upper figure are new punctures. The purple point on the right-hand side of the lower figure, where the two disks are attached, is *not* a puncture.

We now turn to clusters of punctured disks and describe their *admissible degenerations*. Fix a tuple $\bar{\mathcal{L}} = (\bar{L}_0, \dots, \bar{L}_d)$ of Lagrangians. A family Σ^t of $\bar{\mathcal{L}}$-decorated clusters of punctured disks can converge to a decorated degenerate cluster Σ^∞ of punctured disks (which, strictly speaking, by our definitions, might not be a genuine cluster of punctured disks). The degeneration of Σ^t into Σ^∞ can be of several types. The first type is when one (or more) of the punctured disks in the clusters Σ^t degenerates in the way described earlier. Depending on the decoration $\bar{\mathcal{L}}$, one of the two variants mentioned above, or both, can occur. There is one slight exception to this rule. Namely, in both of the variants described above we view the degeneration as admissible only if each of the two punctured disks formed by the splitting contains at least three distinguished points. Here, by a distinguished point we mean either a puncture, a point attached to a tree, or (in case of the second variant) the point of attachment to the other punctured disk in the degenerate cluster.

Below we will describe another four types of admissible degenerations. Before we go into this, a quick remark about the decorations of the limit Σ^∞ is in order.

Our conventions for decorations require the cluster to have an exit point (according to which we label the first entry in the decoration). However, the first variant of the degenerations described above yields two punctured disks S' and S'', where one of them has a (new) exit point and the other has a (new) entry point. The apparent problem is that one of these punctured disks might not have any exit point, hence there might be an ambiguity regarding the order in which we write its decoration. However, this ambiguity is fixed if we use the following conventions. The limit Σ^∞ is divided into two components: the one that contains S' and the one that contains S''. The decorations are uniquely defined once we require that the exit point of the limit Σ^∞ corresponds to the S'' part. A similar thing applies also to clusters in which one of the punctured disks degenerates according to the second variant described earlier.

We now proceed to describe four additional types of admissible degenerations.

Splitting within trees. Apart from degeneration of punctured disks in a cluster, there are several other types of degeneration that can occur within a family Σ^t of clusters of punctured disks. Part of these has to do with degeneration of the trees Σ^t_{tr} of Σ^t, and another part is related to how these trees are attached to Σ^t_S.

Shrinking of edges to a point. The first type of degeneration within trees is when an interior edge in one of the trees of Σ^t_{tr} shrinks to a point (this means that also its interval labeling and parametrization shrink to a point and a constant, respectively). The limit tree will now have one vertex less and will inherit from Σ^t_{tr} all the other structures (such as the labeling of the other edges, the class of planar embedding, etc.). See Figure 3.5. Note, however, that the limit tree will not be admissible (e.g., it might have vertices of valency 4, or a leaf of type A that becomes identified with an interior vertex).

Edge breaking. Another type of degeneration is when one of the interior edges e in a family of trees $T^t \subset \Sigma^t$ becomes of infinite length. We view the limit of the T^t's as a broken tree consisting of two trees T^∞_1 and T^∞_2. These two trees are obtained from T^∞ as follows. We delete the (interior of the) edge e from T^t and obtain two connected components: the part of T^t that appears "before" the edge e and the part that appears "after" that edge (here "before" and "after" refer to the orientation on T^t). Denote by p_1 the end-vertex of the first component (i.e., the entry vertex to the edge e) and by p_2 the new entry vertex of the second component (which corresponds to the exit vertex of e). We now take the first component and add to it a new edge e_1 emanating from p_1. The result is the tree T^∞_1. The other vertex q_1 of e_1 will now be a type-E leaf of T^∞_1 and we regard it as an exit leaf. We label the edge e_1 by $[0,\infty)$, and the rest of the edges are labeled by the limiting labels of T^t as $t \to \infty$. The definition of T^∞_2 is similar, only that now we add a new edge e_2 to the second component (i.e.,

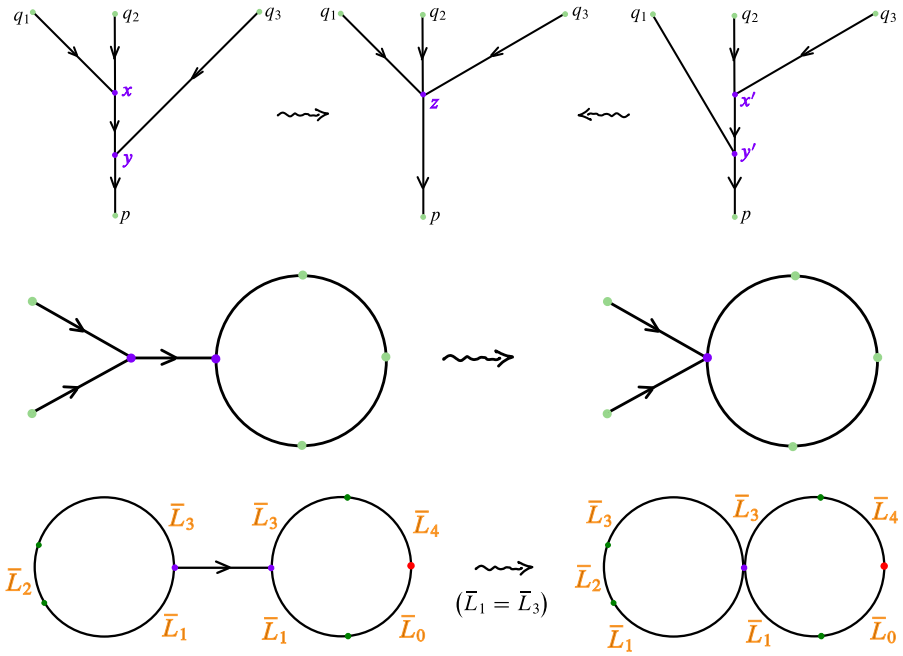


Figure 3.5. Examples of degenerations where an edge of a tree shrinks to a point, leading to a limit (inadmissible) tree, possibly being part of a degenerate cluster.

the one coming "after" the deleted edge e) attached at p_2. The resulting tree is T_2^∞. The new vertex (which is the entry to e_2) will be a type-E leaf of T_2^∞. The edge e_2 is labeled by $(-\infty, 0]$ and the rest of the edges are labeled by the limiting labels of T^t as $t \to \infty$. We refer to T^∞ as a "broken" tree with components T_1^∞ and T_2^∞. See Figure 3.6.

We add the following restriction on edge-breaking degenerations. A degeneration as described above is considered admissible only if neither T_1^∞ nor T_2^∞ is a tree with two vertices both of which are type-E leaves, connected by one edge. All other edge-breaking degenerations are considered admissible.

Collision of type-A leaves. The last type of admissible degeneration is when two type-A leaves (belonging to two different trees) that lie on the boundary of the same punctured disk $S_k \subset \Sigma_S$ collide. This means that two trees $T', T'' \subset \Sigma_{\mathrm{tr}}$ are grafted (or joined) at two of their type-A exit leaves. See Figure 3.7.

Remark 3.11. (a) The boundary of the compactification of the space $\mathrm{Clus}(\overline{\mathcal{L}})$ can be described by the types of degeneration described above. The points of the top-dimensional stratum of the boundary correspond to precisely one such degeneration. Of course, several instances of degeneration can occur simultaneously, but these

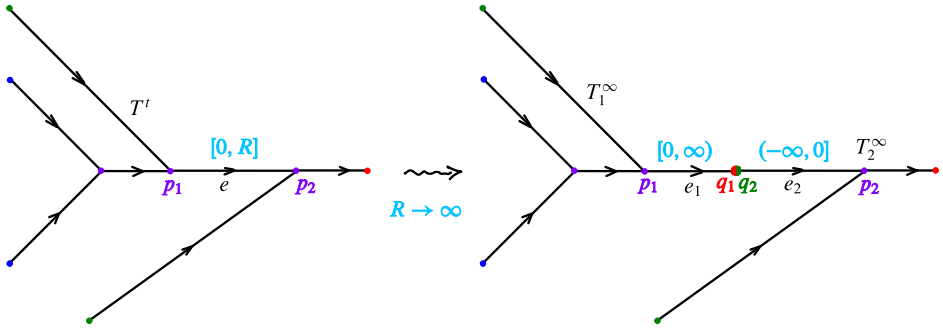


Figure 3.6. Breaking along an edge of a tree, leading to a broken tree with two components.

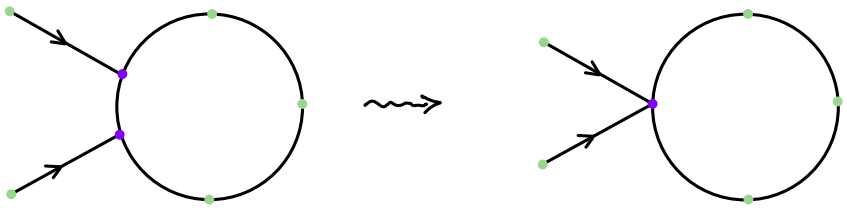


Figure 3.7. Collision of points on ∂S_k to which two different trees are attached.

instances correspond to the lower-dimensional strata of the boundary of $\mathrm{Clus}(\overline{\mathcal{L}})$. In particular, within generic 1-dimensional families of clusters of punctured disks, only one degeneration can occur at a given time.

(b) The converse to "splitting and degeneration" goes by the name gluing. Every degenerate configuration among the ones described above can be obtained as a limit of a family of clusters of punctured disks.

(c) Some of the limit configurations described above can occur as a result of *two different* degenerations. For example, collision of two type-A leaves of trees (along an arc of one of the punctured disks) leads to a configuration which is also the limit of another family of clusters in which one type-A leaf of a tree shrinks to a point. See Figure 3.7 versus the middle part of Figure 3.5.

Similarly, the second variant of splitting within a punctured disk (which leads to two punctured disks connected at a "non-puncture" point along their boundaries) also occurs as a limit of clusters in which two punctured disks are connected by a tree with one edge, and that edge shrinks to a point. See the lower part of Figure 3.4 versus the lower part of Figure 3.5.

The same thing happens with shrinking of interior edges of trees. Namely, each of the (inadmissible) trees that occur after the shrinking of interior edges appears as

a limit of a different family of trees in which an edge shrank to a point. See the left- and right-hand sides of the upper part of Figure 3.5.

The fact that some limit configurations appear in pairs on the total boundary of the compactification of the Clus spaces is important for showing that the μ_d-operations in $\mathcal{F}uk(\overline{\mathcal{X}})$, as will be defined later, satisfy the A_∞-identities. Indeed, when one considers 1-parametric families of clusters of Floer polygons, with fixed entries and exit, some of the boundary points (that correspond to degeneration of the underlying clusters of punctured disks as described above) will appear in pairs and thus can be regarded as "interior" points inside extended families of clusters.

(d) Recall that the boundary of the partial compactification of $\mathrm{Clus}(\overline{\mathcal{L}})$ contains only admissible degenerate clusters. In particular, in the case of splitting of a punctured disk we require that the number of distinguished points on each component is at least 3. Of course, a splitting in which one of the disks has only two distinguished points can occur. The reason we do not add such a configuration to the boundary $\mathrm{Clus}(\overline{\mathcal{L}})$ is that disks with two marked points have a non-compact 1-dimensional group of automorphisms (isomorphic to $\mathbb{R}$). This is referred to in the literature as an unstable marked curve. The situation with the other inadmissible degenerate clusters, namely a broken tree with one component being a tree with one edge connecting between two type-E leaves, is similar. The latter component has an $\mathbb{R}$-action (acting by translation on the identifications between $(-\infty, \infty)$ and the edge of this tree).

In practice, not including these unstable configurations to the boundary of $\mathrm{Clus}(\overline{\mathcal{L}})$ will not cause any problems in showing that the μ_d-operations satisfy the A_∞-identities. The standard way to go about it in Floer and Morse theory is to compactify the space of clusters of *Floer polygons* in such a way that degenerations that correspond to the above unstable configurations are taken into account in the boundary of the latter spaces rather than in $\partial\,\mathrm{Clus}(\overline{\mathcal{L}})$. In terms of the A_∞-identities, these degenerations will contribute the terms in the identities that include μ_1's.

3.3.4 Perturbation data

We assume that Floer data has been chosen for every pair $L_0, L_1 \in \mathcal{X}$, as described at the beginning of Section 3.3.1 on page 109. The perturbation data for a decorated simple cluster (i.e., a cluster consisting of precisely one punctured disk and no trees) S_k is of the same type as in the standard theory, namely it consist of pairs (K, J), where K is a 1-form on S_k with values in the space of compactly supported Hamiltonian functions on the ambient manifold X. This 1-form is assumed to be compatible with the Floer data on each strip-like end of S_k, in the sense that on these ends we have $K \equiv 0$. (Note that we are dealing here with the case of one punctured disk without trees, which means that the decoration $\overline{\mathcal{L}}$ is such that every two consecutive Lagrangians in $\mathcal{L}$ have mutually transverse underlying Lagrangians. Recall also that for pairs of distinct underlying Lagrangians we have already set up the Floer data in advance to

have 0 Hamiltonian terms.) Moreover, we require K to have compact support inside the interior of S_k. The second component of the Floer data is a family of ω-compatible almost complex structures $J = \{J_z\}$ that depends on $z \in S_k$ and coincides on each strip-like end with the almost complex structures chosen for the corresponding Floer data.

We now describe the perturbation data in the case of general decorated clusters of punctured disks Σ. The perturbation data in this case consists of two pieces of data. The first one is a choice of perturbation data (K, J) on each (decorated) punctured disk $S_k \subset \Sigma_S$. The second one is a choice of Morse data on each tree $T \subset \Sigma_{\text{tr}}$ of the cluster. Recall that each such tree T corresponds to an underlying Lagrangian $\overline{L}$ that appears in the decoration $\mathcal{L}$. Recall also that each edge e of T is parametrized by an interval $I_e \subset \mathbb{R}$ (where the intervals for the internal edges are closed of finite length and the ones corresponding to the edges that touch the type-E leaves are semi-infinite). The Morse data for T is a choice of a family $(f_\tau, (\cdot,\cdot)_\tau)$, $\tau \in I_e$, for every edge e of the tree, where for each τ, $f_\tau : \overline{L} \to \mathbb{R}$ is a smooth function and $(\cdot,\cdot)_\tau$ a Riemannian metric on $\overline{L}$. Here $\overline{L}$ is the underlying Lagrangian corresponding to the tree T. Moreover, we require that along the *ends* of the external edges e of T (i.e., the edges connected to the type-E leaves), the pair $(f_\tau, (\cdot,\cdot)_\tau)$ coincides with the Morse data $(f_{\overline{L}}, (\cdot,\cdot)_{\overline{L}})$ associated to $\overline{L}$ that has been fixed in advance. For example, if e is an edge connected to an entry leaf, then $I_e = (-\infty, 0]$, and the requirement is that $(f_\tau, (\cdot,\cdot)_\tau) = (f_{\overline{L}}, (\cdot,\cdot)_{\overline{L}})$ for $\tau \ll 0$. A similar choice of data is made also in case there is an exit leaf (which is the case when $\overline{L}_0 = \overline{L}_d$), only that now the edge e connected to the exit leaf is labeled by $I_e = [0, \infty)$.

There is only one slight exception to the above, namely when $\overline{\mathcal{L}} = (\overline{L}, \overline{L})$. In this case, the whole cluster consists of only one tree (and no punctured disks). This tree has two vertices and one edge e connecting them, which is modeled on the interval $I_e = (-\infty, \infty)$. The choice of Morse data here will be the same Morse data $(f_{\overline{L}}, (\cdot,\cdot))$ chosen in advance for $\overline{L}$, and it is required to be independent of the parameter $\tau \in I_e$.

For every tuple $\overline{\mathcal{L}}$ we make a continuous choice of perturbation data for all the clusters of punctured disks that are parametrized by $\text{Clus}(\overline{\mathcal{L}})$. We denote such a choice by $\mathcal{P}_{\overline{\mathcal{L}}}$ and denote by $\mathcal{P} = \{\mathcal{P}_{\overline{\mathcal{L}}}\}$ the collection of choices $\mathcal{P}_{\overline{\mathcal{L}}}$ made for all tuples $\overline{\mathcal{L}}$ of any length. We refer to $\mathcal{P}$ as a choice of perturbation data.

These choices of $\mathcal{P}$ are subject to being consistent with the splitting, degeneration, and gluing described in Section 3.3.3. This is crucial in order to establish the A_∞-identities among the μ_d-operations that will be described next.

3.3.4.1 The μ_d-operations. We now proceed to the A_∞-operations, taking also filtrations into account. The definition of μ_d, $d \geq 2$, is based on clusters of Floer polygons, which we now describe.

Let $\overline{\mathcal{L}} = (\overline{L}_0, \ldots, \overline{L}_d)$ be a tuple of Lagrangians in $\overline{X}$. An $\overline{\mathcal{L}}$-decorated cluster of Floer polygons is a map $u : \Sigma \to X$ whose domain is an $\overline{\mathcal{L}}$-decorated cluster of punc-

tured disks Σ. The restriction $u|_{S_k}$ of u to any of the punctured disks $S_k \subset \Sigma_S$ is a Floer polygon, exactly as in the standard theory of Fukaya categories [53, Chapter II, Section 9]. Namely, $u|_{S_k}$ satisfies the (generalized) Floer equation associated to the perturbation data that $\mathcal{P}_{\bar{\mathcal{L}}}$ assigns to S_k. The map $u|_{S_k}$ satisfies Lagrangian boundary conditions prescribed by the decoration. The punctures of S_k are sent by u to intersection points of pairs of Lagrangians, as prescribed by the decoration. In addition, we assume that the energy $E(u|_{S_k})$ of $u|_{S_k}$ is finite.

Next, the restriction $u|_T$ of u to any of the trees $T \subset \Sigma_{\mathrm{tr}}$ should satisfy the *negative* gradient equations corresponding to the Morse data specified along the edges of the trees. The interval labeling I_e on the edges e of T and the identifications $\sigma_e^T : e \to I_e$ are used in order to endow each interval with a "time parameter" for the negative gradient trajectories. Finally, the type-E leaves of each tree T from $\mathcal{T}$ are mapped by u to critical points of the functions $f_{\bar{L}}$, where $\bar{L}$ is the Lagrangian decorating the tree T.

Given the choices of Floer and perturbation data, the definition of the μ_d-operations, $d \geq 2$, is now done by counting decorated clusters of Floer polygons with specified boundary conditions and given entry and exit points. More specifically, let $\mathcal{L} = (L_0, \ldots, L_d)$, $d \geq 1$, be a tuple of Lagrangians from $\mathcal{X}$, and denote by $\bar{\mathcal{L}} = (\bar{L}_0, \ldots, \bar{L}_d)$ the corresponding tuple of underlying Lagrangians. Define

$$\begin{aligned} &\mu_d : \mathrm{CF}(L_0, L_1) \otimes \cdots \otimes \mathrm{CF}(L_{d-1}, L_d) \to \mathrm{CF}(L_0, L_d), \\ &\mu_d(x_1, \ldots, x_d) := \sum_y \#\mathcal{M}(x_1, \ldots, x_d, y; \mathcal{P})\, y. \end{aligned} \tag{3.6}$$

Here we have abbreviated $\mathrm{CF}(L', L'') := \mathrm{CF}(L', L''; \mathcal{D}_{L',L''})$ for any $L', L'' \in \mathcal{X}$. The sum in the second line of (3.6) goes as follows: y runs over all the generators of $\mathrm{CF}(L_0, L_d)$ of appropriate degree, and $\#\mathcal{M}(x_1, \ldots, x_d, y; \mathcal{P})$ stands for the count (with values in $\mathbb{Z}_2$, or under additional assumptions in $\mathbf{k}$) of elements in the 0-dimensional component of the space $\mathcal{M}(x_1, \ldots, x_d, y; \mathcal{P})$ of $\bar{\mathcal{L}}$-decorated clusters of Floer polygons with entry points $x_1, \ldots, x_d$ and exit point y.

We denote by $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$ the collection of objects $\mathcal{X}$ together with the multilinear operations μ_d, $d \geq 1$, associated to the Floer and perturbation data $\mathcal{P}$. Notice that the perturbation data depends only on the geometric part of the marked Lagrangians, namely $\bar{\mathcal{X}}$.

Theorem 3.12. *For every finite collection of Lagrangians $\bar{\mathcal{X}}$ satisfying the conditions from the beginning of Section 3.2.1, there exists a (non-empty) space $\mathcal{B}(\bar{\mathcal{X}})$ of regular Floer and perturbation data, of the types described above, such that for every $\mathcal{P} \in \mathcal{B}(\bar{\mathcal{X}})$ the following holds:*

(i) *$\mathcal{F}uk(\mathcal{X}; \mathcal{P})$, with the above μ_d-operations, $d \geq 1$, is a strictly unital A_∞-category. Moreover, with the filtrations defined in Section 3.3.1, this A_∞-category is genuinely filtered.*

(ii) *If one forgets the filtrations, then $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ as defined above is quasi-equivalent to the subcategory of the standard Fukaya category (e.g., as defined in [53]) whose collection of objects is $\mathcal{X}$. This quasi-equivalence can be assumed to be the identity map on the set of objects $\mathcal{X}$.*

Moreover, there exist filtered A_∞-functors

$$\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0} : \mathcal{F}uk(\mathcal{X};\mathcal{P}_0) \to \mathcal{F}uk(\mathcal{X};\mathcal{P}_1),$$

defined for every $\mathcal{P}_0, \mathcal{P}_1 \in \mathcal{B}(\bar{\mathcal{X}})$, with the following properties:

(1) *$\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ is A_1-unital (see Section 3.3.6 for the definition).*

(2) *$\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ is a filtered quasi-equivalence.*

(3) *The action of $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ on objects is the identity map. (Recall that all the categories $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ have the same set of objects $\mathcal{X}$.)*

(4) *For every $L', L'' \in \mathcal{X}$, the maps $(\mathcal{F}_1^{\mathcal{P}_1,\mathcal{P}_0})_* : \mathrm{HF}(L', L'';\mathcal{P}_0) \to \mathrm{HF}(L', L'';\mathcal{P}_1)$ induced by the first order components of $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ are the canonical continuation isomorphisms in Floer theory.*

(5) *$\mathcal{F}^{\mathcal{P},\mathcal{P}} = \mathbb{1}$.*

(6) *The composition $\mathcal{F}^{\mathcal{P}_2,\mathcal{P}_1} \circ \mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ is isomorphic to $\mathcal{F}^{\mathcal{P}_2,\mathcal{P}_0}$ in the category $H^0(\mathit{ffun}(\mathcal{F}uk(\mathcal{X};\mathcal{P}_0), \mathcal{F}uk(\mathcal{X};\mathcal{P}_2)))_0$. Here, $\mathit{ffun}(\mathcal{F}uk(\mathcal{X};\mathcal{P}_0), \mathcal{F}uk(\mathcal{X};\mathcal{P}_2))$ is the category of filtered A_∞-functors from $\mathcal{F}uk(\mathcal{X};\mathcal{P}_0)$ to $\mathcal{F}uk(\mathcal{X};\mathcal{P}_2)$, $H^0(\mathit{ffun}(\cdots))$ is the persistence homological category of ffun in cohomological degree 0, and $H^0(\mathit{ffun}(\cdots))_0$ is its 0-level persistence subcategory. In other words, there exists an A_∞-natural transformation $T^{\mathcal{P}_2,\mathcal{P}_1,\mathcal{P}_0} : \mathcal{F}^{\mathcal{P}_2,\mathcal{P}_1} \circ \mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0} \to \mathcal{F}^{\mathcal{P}_2,\mathcal{P}_0}$ that preserves filtrations and is an isomorphism in the homological persistence category of filtered functors. Note that, in particular, this implies that $\mathcal{F}^{\mathcal{P}_0,\mathcal{P}_1} \circ \mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ is isomorphic to $\mathbb{1}_{\mathcal{F}uk(\mathcal{X};\mathcal{P}_0)}$ in the respective homological persistence category.*

Furthermore, the choice of the assignment $\bar{\mathcal{X}} \mapsto \mathcal{B}(\bar{\mathcal{X}})$ can be assumed to have the following property: if $\bar{\mathcal{X}}'$ is another finite collection of Lagrangians with $\bar{\mathcal{X}}' \supset \bar{\mathcal{X}}$ that (similarly to $\bar{\mathcal{X}}$) satisfies the conditions from the beginning of Section 3.2.1, then $\mathcal{B}(\bar{\mathcal{X}}')|_{\bar{\mathcal{X}}} \subset \mathcal{B}(\bar{\mathcal{X}})$. Here, $\mathcal{B}(\bar{\mathcal{X}}')|_{\bar{\mathcal{X}}}$ stands for the restriction of the perturbation data from $\mathcal{B}(\bar{\mathcal{X}}')$ to the spaces of clusters of punctured disks decorated by the elements of $\bar{\mathcal{X}}$.

Remark 3.13. The system of functors $\{\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}\}_{\mathcal{P}_1,\mathcal{P}_0 \in \mathcal{B}(\bar{\mathcal{X}})}$ and the natural transformations mentioned in Theorem 3.12 depend on a variety of choices and hence are not canonical in the strict sense of the word. The extent to which these structures are canonical will be briefly discussed later in Remark 3.24. We will refer to this system of functors as a *weakly coherent system* of comparison functors to emphasize that our construction does not produce canonical choices.

We proceed now to the proof of Theorem 3.12. The proof presented below is by no means complete and should be viewed as an outline only. We have left out quite a few technical details, especially concerning the analysis underlying the proof. However, these parts of the proof follow from rather standard and well-established ingredients in the analysis of Floer theory and Fukaya categories. As mentioned earlier, a more general approach to genuinely filtered Fukaya categories is worked out in a forthcoming paper by Ambrosioni [1], which will also contain a detailed proof of the construction.

3.3.5 Proof of Theorem 3.12, part 1

We will concentrate here on the second part of point (i) of the theorem (namely that $\mathcal{F}uk(\mathcal{X})$ can be made *genuinely filtered* for appropriate choices of Floer and perturbation data). The first part of point (i) and point (ii) are rather well known and have been addressed in the literature with various levels of rigor. The proofs of the statements concerning the functors $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ will be outlined in Section 3.3.6 below.

Throughout the proof we will sometimes abbreviate $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ as $\mathcal{F}uk(\mathcal{X})$ in case $\mathcal{P}$ is clear from the context.

Fix a tuple of Lagrangians $\bar{\mathcal{L}} = (\bar{L}_0, \dots, \bar{L}_d)$ and assume for simplicity that $\bar{L}_i \neq \bar{L}_j$ for every $i \neq j$. Fix a tuple of intersection points $x_1, \dots, x_d, y$, where $x_i \in \bar{L}_{i-1} \cap \bar{L}_i$ and $y \in \bar{L}_0 \cap \bar{L}_d$. Let $u \in \mathcal{M}(x_1, \dots, x_d, y; \mathcal{P})$ be an $\bar{\mathcal{L}}$-decorated cluster of Floer polygons. By our simplifying assumptions, the domain of u must be a simple cluster, namely just one punctured disk S_{d+1}. (See page 114 and Remark 3.10.) Thus $u : S_{d+1} \to X$ is a Floer polygon that sends the punctures of S_{d+1} to the points $x_1, \dots, x_d, y$. For such a map u, denote by $A(u) = \int_{S_{d+1}} u^*\omega$ the symplectic area of u. We have

$$A(u) = \sum_{i=1}^{d} \mathcal{A}(x_i) - \mathcal{A}(y).$$

In order to prove that the μ_d-operation preserves filtrations, we need to show that for all u as above we have $A(u) \geq 0$. (In fact, we need to show the latter inequality holds for all clusters of Floer polygons u, not only the simple ones. However, as we will see below, the main difficulty is for punctured disks, and the generalization to more general clusters is straightforward.)

Before we go into the proof of the latter statement, let us explain the difficulties underlying it. Obviously, $A(u) \geq 0$ if we choose the perturbation data (K, J) with $K \equiv 0$, since then every Floer polygon will be J-holomorphic and hence of strictly positive area. However, for a variety of reasons it seems better to allow for non-trivial 1-forms K in the perturbation data, so we will not assume $K \equiv 0$. One of the reasons for allowing non-trivial perturbations is that it is easier to establish transversality for the spaces of Floer clusters with this extra parameter at hand. Another reason is that if

one hopes to generalize the present approach to cases when not all pairs of Lagrangians in $\bar{\mathcal{X}}$ intersect transversely then Hamiltonian perturbations would definitely be needed. Other reasons have to do with compatibility of $\mathcal{F}uk(\mathcal{X})$ with other structures, such as the Floer (or symplectic) homology of the ambient manifold and maps relating $\mathcal{F}uk(\mathcal{X})$ to these invariants. These structures usually require Hamiltonian perturbations.

We may next try to show that $A(u) \geq 0$ once we take the 1-forms K in the perturbation data to be small enough. At first sight, this seems to work using a compactness argument as follows. Indeed, if this were not the case, then we would have a sequence of perturbation forms $K^{(l)}$ and a sequence $u_l : S^{(l)}_{d+1} \to X$ of corresponding Floer polygons whose punctures go to a fixed set of intersection points $x_1, \ldots, x_d, y$, such that $K^{(l)} \xrightarrow{C^1} 0$ but $A(u_l) \leq 0$ for every l. (We may assume that all the u_l's run between the *same* set of intersection points $x_1, \ldots, x_d, y$, because by our assumptions there is only a finite number of possible intersection points associated to the tuple $\bar{\mathcal{L}}$.) By compactness, passing to a subsequence of the u_l's we would then obtain a (possibly broken) limit polygon u which is non-constant and genuinely J-holomorphic, yet with $A(u) \leq 0$. A contradiction. Since, by assumption, the number of possible $d+1$ tuples $\bar{\mathcal{L}}$ is finite, it follows that, if we take the perturbation data small enough, then for all tuples $\bar{\mathcal{L}}$ of length $d+1$ with the properties from the beginning of the proof, and all $\bar{\mathcal{L}}$-decorated Floer polygons u, we have $A(u) > 0$. This easily extends also to decorated clusters of Floer polygons that are not necessarily simple, as well as to all decorations $\bar{\mathcal{L}}$ of *fixed* length $d+1$.

The problem with this argument is that, without further elaboration, it might create difficulties with obtaining a consistent choice of perturbation data $\mathcal{P}$. To explain this difficulty, let us rephrase the previous paragraph in more quantitative terms. From now on we will use the following more detailed notation. We denote the restriction of the perturbation data $\mathcal{P}$ to the space of all $(d+1)$-punctured disks by $\mathcal{P}_{d+1}$, and write $(K(\mathcal{P}_{d+1}), J(\mathcal{P}_{d+1}))$ for the two components of $\mathcal{P}_{d+1}$. Given an $\bar{\mathcal{L}}$-decorated punctured disk S_{d+1}, we denote by $(K(\mathcal{P}, S_{d+1}), J(\mathcal{P}, S_{d+1}))$ the restriction of $\mathcal{P}$ to S_{d+1}.

The previous argument shows that there exist numbers $\epsilon_{d+1} > 0$, $d \geq 2$, such that if the perturbation forms $K(\mathcal{P}_{d+1})$ satisfy $\|K(\mathcal{P}_{d+1})\| \leq \epsilon_{d+1}$ then for all tuples $\mathcal{L}$ of length $d+1$ and all $\bar{\mathcal{L}}$-decorated Floer polygons u, we have $A(u) > 0$. Here $\|-\|$ is a suitable norm on the space of all perturbation 1-forms K_{d+1} (defined on the space of all possible $(d+1)$-punctured disks S_{d+1}). The value of $\|K_{d+1}\|$ involves the values of K_{d+1} and its first derivatives, both in the domain direction as well as in the direction of the manifold X (recall that the forms K_{d+1} have values in the space of compactly supported functions on X).

The problem that arises with the approach used so far has to do with the consistency of $\mathcal{P}$ with respect to gluing/splitting. A standard way to construct consistent perturbation data is to construct $K(\mathcal{P}_{d+1})$ (and the almost complex structures) by

induction on d and to ensure that at each induction step the newly defined perturbation data is consistent with the data that have already been defined at earlier stages, with respect to gluing/splitting. Assume that $K(\mathcal{P}_{m+1})$ has already been defined for all $m \leq d_0$ in such a way that $\|K(\mathcal{P}_{m+1})\| \leq \epsilon_{m+1}$ for all $m \leq d_0$. Consider now $d = d_0 + 1$. Punctured disks of the type S_{d+1} can split into two punctured disks of the type $S_{d'+1}$ and $S_{d''+1}$, with $d' + d'' = d + 1$ (and $d', d'' \geq 2$). See Figure 3.8.

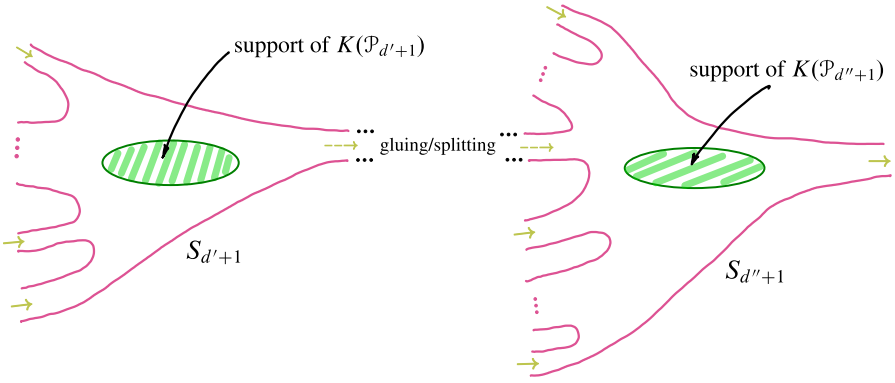


Figure 3.8. Gluing/splitting of punctured disks. Two punctured disks of the types $S_{d'+1}$ and $S_{d''+1}$, together with their perturbation 1-forms, are glued into a punctured disk of the type S_{d+1}, where $d + 1 = d' + d''$.

If it so happens that

$$\epsilon_{d+1} < \|K(\mathcal{P}_{d'+1})\| + \|K(\mathcal{P}_{d''+1})\|,$$

then the induction step will produce perturbation forms $K(\mathcal{P}_{d+1})$ that might not satisfy the condition $\|K(\mathcal{P}_{d+1})\| \leq \varepsilon_{d+1}$, which is required in order to have $A(u) > 0$ for all Floer polygons u. In other words, the construction described above does not go through at the induction step. Of course, the above argument shows that for every fixed d, we can choose consistent perturbation data such that all the Floer polygons involving no more than $d + 1$ intersection points have positive symplectic area. One could then easily modify this argument to handle more general clusters and consequently show that all the μ_k-operations in $\mathcal{F}uk(\mathcal{X})$ with $k \leq d$ will preserve filtrations. However, without any additional arguments, the above fails to prove that there is even one consistent perturbation data that will turn the μ_k-operations into filtration-preserving maps *for all* k. We will now refine the above arguments, showing how to achieve our goal by restricting further the type of perturbation data.

Recall that for all $\mathcal{L}$-decorated Floer polygons $u : S_{d+1} \to X$ we have the following energy-area identity:

$$E(u) = A(u) - \int_{S_{d+1}} R_{K(\mathcal{P}, S_{d+1})} \circ u, \tag{3.7}$$

where $E(u)$ is the energy of u and $R_{K(\mathcal{P},S_{d+1})}$ is the curvature associated to the 1-form $K(\mathcal{P}, S_{d+1})$. Recall that $R_{K(\mathcal{P},S_{d+1})}$ is a 2-form on S_{d+1} with values in the space of compactly supported functions on X. The expression $R_{K(\mathcal{P},S_{d+1})} \circ u$ is a real-valued 2-form on S_{d+1} obtained by composing the functions prescribed by the values of $R_{K(\mathcal{P},S_{d+1})}$ with the map u. We refer the reader to [53, Chapter II, Section (8g)] for more details on the definition of $R_{K(\mathcal{P},S_{d+1})}$ and the identity (3.7).

An important point about the curvature form $R_{K(\mathcal{P},S_{d+1})}$ is that it can be made arbitrarily small by choosing the perturbation form $K(\mathcal{P}, S_{d+1})$ to be small enough in an appropriate C^1-norm $\|-\|$. As alluded to above, $\|K(\mathcal{P}, S_{d+1})\|$ involves the sum of the L^1-norms of $K(\mathcal{P}, S_{d+1})$ and its derivatives, both in the direction of S_{d+1} and in the direction of the manifold X. (This norm can be viewed as a version of the Sobolev norm on $W^{1,1}$.) We therefore have

$$\left| \int_{S_{d+1}} R_{K(\mathcal{P},S_{d+1})} \circ v \right| \leq C_{d+1} \|K(\mathcal{P}, S_{d+1})\|$$

for all maps $v : S_{d+1} \to X$, where the constant C_{d+1} depends only on d (and not on the specific surface S_{d+1} or any other parameter involved in the integrand). Define

$$\begin{aligned} \|K(\mathcal{P}_{d+1})\| &:= \sup \|K(\mathcal{P}, S_{d+1})\|, \\ \nu(\mathcal{P}_{d+1}) &:= C_{d+1} \|K(\mathcal{P}_{d+1})\|, \end{aligned} \tag{3.8}$$

where the supremum in the first formula is taken over all $\overline{\mathcal{L}}$-decorated punctured disks S_{d+1}. Note that it is possible to choose the perturbation forms $K(\mathcal{P}, S_{d+1})$ such that they are all compactly supported inside the interior of each punctured disk S_{d+1}, and moreover we can control these supports so that the support of the entire family $K(\mathcal{P}_{d+1})$ is compact. (This poses no problems to having consistency with respect to gluing/splitting.) This implies that the supremum that appears in the definition of $\|K(\mathcal{P}_{d+1})\|$ in (3.8) can be assumed to be finite, and moreover can be made arbitrarily small by an appropriate choice of $K(\mathcal{P}_{d+1})$.

With the above notation we now have

$$A(u) \geq E(u) - \nu(\mathcal{P}_{d+1}) \tag{3.9}$$

for all $\overline{\mathcal{L}}$-decorated Floer polygons u.

We now add further restrictions to the perturbation data $\mathcal{P}$. Denote by $\overline{\mathcal{X}} = \{\overline{L} \mid L \in \mathcal{X}\}$ the set of all underlying Lagrangians from the collection $\mathcal{X}$. Denote by $I = \{p_1, \ldots, p_N\}$ the set of all intersection points between any two distinct Lagrangians $\overline{L}, \overline{L}' \in \overline{\mathcal{X}}$. By our assumption on $\overline{\mathcal{X}}$, I is a finite set. Moreover, every $p \in I$ corresponds to precisely one pair of distinct Lagrangians $\overline{L}, \overline{L}' \in \overline{\mathcal{X}}$ whose intersection contains p.

Denote by $B^{2n}(R) \subset \mathbb{R}^{2n}$ the closed $2n$-dimensional ball endowed with its standard symplectic structure, where $2n = \dim_{\mathbb{R}} X$. We claim that there exists a symplectic embedding of a disjoint union of N balls of some radius $R_{\bar{\mathcal{X}}} > 0$ into X,

$$\phi : \bigsqcup_{j=1}^{N} B^{2n}(R_{\bar{\mathcal{X}}}) \to X, \tag{3.10}$$

with the following properties:

- $\phi(0_j) = p_j$ for every j. Here and in what follows, we denote by B_j the jth ball in the disjoint union in (3.10), and by $0_j \in B_j$ the center of that ball.
- If $p_j \in \bar{L}' \cap \bar{L}''$ with $\bar{L}', \bar{L}'' \in \bar{\mathcal{X}}$ distinct, then

$$\begin{aligned} &\text{either } \phi(\mathbb{R}B_j) \subset \bar{L}',\ \phi(i\mathbb{R}B_j) \subset \bar{L}'', \\ &\text{or } \phi(\mathbb{R}B_j) \subset \bar{L}'',\ \phi(i\mathbb{R}B_j) \subset \bar{L}', \end{aligned}$$

  where $\mathbb{R}B_j := B_j \cap (\mathbb{R}^n \times \{0\})$ is the "real" part of the ball B_j and $i\mathbb{R}B_j := B_j \cap (\{0\} \times \mathbb{R}^n)$ is its "imaginary" part. Moreover, $(\phi|_{B_j})^{-1}(\bar{L}' \cup \bar{L}'') \subset \mathbb{R}B_j \cup i\mathbb{R}B_j$, and $\phi(B_j)$ is disjoint from all the Lagrangians in $\bar{\mathcal{X}} \setminus \{\bar{L}', \bar{L}''\}$.

Clearly there exist an $R_{\bar{\mathcal{X}}} > 0$ and an embedding ϕ as above, and we fix both of them once and for all. Next we fix $0 < R < R_{\bar{\mathcal{X}}}$ and denote by $B_j(R) \subset B_j$ the smaller ball of radius R. Fix another parameter $0 < r \leq R$ and impose the following further restrictions on the perturbation data $\mathcal{P}_{d+1} = (K(\mathcal{P}_{d+1}), J(\mathcal{P}_{d+1}))$ for all d:

(1) $K|_{\phi(B_j(R))} \equiv 0$ for every j, i.e., the values of the 1-form K (which are compactly supported functions on X) vanish over the image of the restriction of ϕ to the smaller balls of radius R.

(2) $J_z|_{\phi(B_j(R))} = \phi_*(J_{\text{std}})$ for every j and every $z \in S_{d+1}$, $d \geq 2$. Here J_{std} is the standard complex structure on $B_j(R)$.

Note that the above additional restrictions on (K, J) do not contradict any of our previous assumptions on (K, J), and if we temporarily ignore the size of the perturbation forms K, then the above also do not pose any problems to the consistency of $\mathcal{P}$ with respect to gluing/splitting.

We now claim that there is a *consistent* choice of perturbation data $\mathcal{P}$ of the type described above such that for all Floer polygons u we have $A(u) > 0$.

To prove this, consider an $\bar{\mathcal{L}}$-decorated Floer polygon $u : S_{d+1} \to X$ associated with our perturbation data. Denote by $x_1, \ldots, x_d, y$ the intersection points to which the punctures of S_{d+1} are mapped by u. By construction, u is genuinely holomorphic over image(ϕ) with respect to an almost complex structure that is diffeomorphic along that region to J_{std} via ϕ. Denote by $B_{j_1}, \ldots, B_{j_d}, B_{j_{d+1}}$ the balls corresponding to the intersection points $x_1, \ldots, x_d, y$ according to the construction made earlier. Applying a version of the monotonicity lemma (or, alternatively, a version of the

Lelong inequality) to u over $\phi(B_{j_k})$, $k = 1, \dots, d+1$, we obtain that there exists a constant C that does not depend on d (or on u or on any other parameter from the perturbation data) such that

$$E(u) \geq (d+1)CR^2. \tag{3.11}$$

Putting this together with (3.9) we obtain

$$A(u) \geq (d+1)CR^2 - \nu(\mathcal{P}_{d+1}). \tag{3.12}$$

Obviously, if we choose perturbation data $\mathcal{P}$ that are small enough such that

$$\nu(\mathcal{P}_{d+1}) < (d+1)Cr^2, \tag{3.13}$$

then $A(u) > 0$ for all Floer polygons u. (Recall that we fixed the parameter r such that $r \leq R$.) The main thing that needs to be verified now is that the condition (3.13) still enables a choice of perturbation data that are consistent with gluing/splitting. We address this point next.

We will choose the perturbation data $\mathcal{P}$ such that for all d we have $\nu(\mathcal{P}_{d+1}) < \alpha_{d+1}(d+1)Cr^2$ for some $0 < \alpha_{d+1} \leq 1$. Recall that punctured disks of the type S_{d+1} can split into two punctured disks of the type $S_{d'+1}$ and $S_{d''+1}$, with $d' + d'' = d+1$ and $d', d'' \geq 2$. (Of course, splittings into more than two punctured disks are also possible; however, for the sake of obtaining a filtered A_∞-category, the top strata of the boundary of the space of disks matter, and these correspond to splitting into two punctured disks only.) And vice versa, by gluing two punctured disks of the type $S_{d'+1}$ and $S_{d''+1}$, we obtain punctured disks of the type $S_{d'+d''}$. Therefore, in order to make it possible to construct the perturbation data $\mathcal{P}_{d+1}$ by induction on d as indicated earlier, and to make them consistent with respect to gluing/splitting, we need to find a sequence of numbers α_{d+1}, $d \geq 2$, that satisfy the following set of inequalities:

$$\begin{aligned} &\alpha_{d'+1}(d'+1)Cr^2 + \alpha_{d''+1}(d''+1)Cr^2 \leq \alpha_{d'+d''}(d'+d'')Cr^2 \quad \forall\, d', d'' \geq 2, \\ &0 < \alpha_{d+1} \leq 1 \quad \forall\, d \geq 2. \end{aligned} \tag{3.14}$$

If such a sequence of numbers α_k does exist then we simply construct the perturbation data $\mathcal{P}_{d+1}$ by induction on d, where at each induction step we require that $\nu(\mathcal{P}_{d+1}) < \alpha_{d+1}(d+1)Cr^2$. The inequalities (3.14) will then ensure that the induction step goes through without posing problems to consistency with respect to gluing/splitting.

It remains to show that the inequalities in (3.14) admit solutions. Setting $\beta_k := k\alpha_k$ for every $k \geq 3$, the set of inequalities (3.14) can be simplified to

$$\begin{aligned} &\beta_{d'+1} + \beta_{d''+1} \leq \beta_{d'+d''} \quad \forall\, d', d'' \geq 2, \\ &0 < \beta_k \leq k \quad \forall\, k \geq 3. \end{aligned} \tag{3.15}$$

It is easy to see that this set of inequalities does have solutions. For example, for the sequence $\beta_k := (3k-6)B, k \geq 3$, where $B \leq \frac{1}{3}$, the first inequality in (3.15) becomes an equality and the second inequality is satisfied. One can also find sequences for which all inequalities in (3.15) become strict. This can be done as follows. Let β'_k be any sequence for which the first line in (3.15) is an equality and the second inequality holds (e.g., the preceding sequence $(3k-6)B$ with $B \leq \frac{1}{3}$). Let $\sigma : [3, \infty) \to \mathbb{R}$ be a strictly increasing function with $0 < \sigma(x) < 1$ for every x. Define now $\beta_k := \sigma(k)\beta'_k$ for all $k \geq 3$. A straightforward calculation shows that for this sequence all the inequalities in (3.15) become strict.

To conclude the proof regarding the preservation of filtrations of the μ_k-operations we need to address also the case of non-simple clusters of Floer polygons. The argument in this case is essentially the same, and below we will only outline it in the case of a cluster consisting of at most one Floer polygon with possibly several trees attached to its boundary.

Consider first a tuple $\bar{\mathcal{L}} = (\bar{L}_0, \ldots, \bar{L}_d)$ in which not all the Lagrangians coincide. Fix a tuple $x_1, \ldots, x_d, y$ of generators, with $x_i \in \mathrm{CF}(L_{i-1}, L_i; \mathcal{D}_{L_{i-1},L_i})$, $y \in \mathrm{CF}(L_0, L_d; \mathcal{D}_{L_0,L_d})$. (Recall that some of these points are intersection points between the respective underlying Lagrangians, and some are critical points of the Morse functions prescribed by the Morse data.) Let Σ be an $\bar{\mathcal{L}}$-decorated cluster of punctured disks consisting of one punctured disk $S_{d'+1}$, where $d' \leq d$, and several trees. Let $u : \Sigma \to X$ be an $\bar{\mathcal{L}}$-decorated cluster of Floer polygons with entry points $x_1, \ldots, x_d$ and exit point y. Denote by $u' := u|_{S_{d'+1}}$ the restriction of u to the underlying punctured disk $S_{d'+1}$ (which is a genuine Floer polygon with boundary conditions prescribed by a sub-tuple of $\bar{\mathcal{L}}$).

Recalling our filtration conventions for $\mathrm{CF}(L', L'')$ in case $\bar{L}' = \bar{L}''$ (see at the end of Section 3.3.1), a simple calculation shows that

$$\sum_{i=1}^{d} \mathcal{A}(x_i) - \mathcal{A}(y) = A(u').$$

Now essentially the same argument as the one carried out earlier shows that the perturbation data can be chosen such that $A(u') \geq 0$.

Let us also consider the case when all the Lagrangians in $\bar{\mathcal{L}}$ coincide. In that case, a cluster of Floer polygons is just a collection of Morse trajectories modeled on a tree (without any actual polygons). A simple calculation shows that we have

$$\mathcal{A}(y) = \sum_{i=1}^{d} \mathcal{A}(x_i),$$

which implies that the filtration is preserved by the operations μ_d, $d \geq 2$, also in the case when the Lagrangians in $\bar{\mathcal{L}}$ all coincide.

As mentioned above, these arguments easily generalize to more complicated clusters, though the notation becomes more involved.

Transversality. We will briefly address now the topic of transversality. In order to show that $\mathcal{F}uk(\mathcal{X})$ is indeed an A_∞-category (i.e., that μ_d-operations satisfy the A_∞-identities, etc.), one needs to choose perturbation data $\mathcal{P}$ that satisfy various regularity properties. This would ensure that the spaces of clusters of Floer polygons involved in the definition of the μ_d-operations are smooth manifolds and have other desirable properties. Establishing the existence of regular perturbation data (or other auxiliary data) usually goes by the name "transversality", and is carried out via analytic techniques that have become standard in Floer theory. The typical result in this context is that the set of regular perturbation data is residual (in particular, dense) inside the space of all consistent perturbation data. However, in our case we work within a much more restricted space of perturbation data, as described above (e.g., specific choices of almost complex structures near the intersection points between pairs of distinct Lagrangians, perturbation forms that are compactly supported, etc.). Formally speaking, one would need to work out the transversality for our choices of perturbation data. While this does not formally follow from the general transversality theorems, it can still be achieved by rather standard arguments. For example, the fact that we restrict the almost complex structures to be constant on certain regions does not pose any problems (for achieving regularity), as long as the images of all the Floer polygons pass through regions of X in which we are allowed to perturb the almost complex structures without any restrictions. The same goes for the perturbation 1-forms K. In a similar vein, the fact that our perturbation forms K must be chosen to be small enough does not pose any transversality problem either. The only ingredient that requires slightly different transversality arguments is the part that uses the combination of Morse trees and Floer polygons. Transversality for spaces of Morse trees, as well as Morse trees mixed with holomorphic curves, has been worked out in various setups; see [15, 16, 32]. While we do not provide here details for these arguments, the case of clusters of Floer polygons follows from the source space description for clusters that appears in [15] and standard regularity arguments, as outlined above. Because all our Lagrangians are exact, the actual regularity arguments are much simpler compared to the ones developed in [15, 16].

The spaces $\mathcal{B}(\bar{\mathcal{X}})$. So far the proof above shows that there exist regular perturbation data $\mathcal{P}$ that turn $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ into a genuinely filtered, strictly unital, A_∞-category.

We now define the space $\mathcal{B}(\bar{\mathcal{X}})$. Recall that $R_{\bar{\mathcal{X}}}$ is the radius of the balls in the embedding ϕ. Given $r \leq R \leq R_{\bar{\mathcal{X}}}$, we denote by $\mathcal{B}(\bar{\mathcal{X}}; R, r)$ the space of consistent, regular perturbation data $\mathcal{P}$, as defined earlier in the proof with the parameters r and R. The significance of the parameter R appears in conditions (1)–(2) on page 129 regarding the perturbation forms and almost complex structures on $\bigcup_j \phi(B_j(R))$. The parameter r plays a role in the size of the perturbations in condition (3.13) on

page 130. To define $\mathcal{B}(\bar{\mathcal{X}})$ we specialize $\mathcal{B}(\bar{\mathcal{X}}; R, r)$ to the case $r = R$, and take the union over all $R \leq R_{\bar{\mathcal{X}}}$. More precisely,

$$\mathcal{B}(\bar{\mathcal{X}}) := \bigcup_{0<R\leq R_{\bar{\mathcal{X}}}} \mathcal{B}(\bar{\mathcal{X}}; R, R).$$

This completes the outline of the proof of the second part of the first statement of Theorem 3.12. ∎

Remark 3.14. (a) The reason for introducing the more general spaces $\mathcal{B}(\bar{\mathcal{X}}; R, r)$ (instead of working with $r = R$ all the time) will become clear in Section 3.3.6 when we prove the results on the system of functors $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$.

(b) If $R' \neq R''$ then neither $\mathcal{B}(\bar{\mathcal{X}}; R', R')$ nor $\mathcal{B}(\bar{\mathcal{X}}; R'', R'')$ is a subspace of the other. In fact, if $R'' \leq R'$ then we have

$$\mathcal{B}(\bar{\mathcal{X}}; R', R') \cap \mathcal{B}(\bar{\mathcal{X}}; R'', R'') = \mathcal{B}(\bar{\mathcal{X}}; R', R'').$$

Moreover, if $R'' \leq R'$ and $r'' \geq r'$, then $\mathcal{B}(\bar{\mathcal{X}}; R', r') \subset \mathcal{B}(\bar{\mathcal{X}}; R'', r'')$.

(c) Obviously, the space $\mathcal{B}(\bar{\mathcal{X}})$ does not only depend on $\bar{\mathcal{X}}$ but also on the choice of the embedding ϕ, and there does not seem to be any preferred choice in this respect.

3.3.6 Proof of Theorem 3.12, part 2: Coherent systems

Here we will follow Seidel's approach to invariance of Fukaya categories based on coherent systems as described in [53, Chapter II, Section (10a)], but with several modifications needed to accommodate the filtered setting.

Before we go on, we briefly explain what is meant by a coherent system of *filtered* A_∞-categories. This is the filtered counterpart of the notion of coherent systems of A_∞-categories introduced in [53, Section (10a), pp. 133–135]. More specifically, let $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$ be a family of filtered A_∞-categories. A coherent system consists of A_∞-functors $\mathcal{F}^{i_1,i_0} : \mathcal{A}^{i_0} \to \mathcal{A}^{i_1}$, defined for all $i_0, i_1 \in \mathcal{I}$, and with $\mathcal{F}^{i,i} = \mathbb{1}_{\mathcal{A}^i}$ for all i, as well as natural transformations $T^{i_2,i_1,i_0} : \mathcal{F}^{i_2,i_1} \circ \mathcal{F}^{i_1,i_0} \to \mathcal{F}^{i_2,i_0}$ for all $i_0, i_1, i_2 \in \mathcal{I}$. The functors $\mathcal{F}^{i,j}$ will be called comparison (or transition) functors. The functors $\mathcal{F}^{i,j}$ and natural transformations $T^{i,j,k}$ are required to satisfy a list of conditions explained in [53, Section (10a), p. 134]. In particular, the comparison functors $\mathcal{F}^{i_1,i_0}$ are all quasi-equivalences and the $T^{i,j,k}$ are quasi-isomorphisms. Turning to the filtered case, we require the following additional conditions to hold. All the functors $\mathcal{F}^{i_1,i_0}$ are required to be filtered (i.e., filtration preserving) and the T^{i_2,i_1,i_0} should be natural transformations of shift-0, i.e.,

$$T^{i_2,i_1,i_0} \in \hom^{\leq 0}_{\mathit{ffun}}(\mathcal{A}^{i_2}, \mathcal{A}^{i_0})(\mathcal{F}^{i_2,i_1} \circ \mathcal{F}^{i_1,i_0}, \mathcal{F}^{i_2,i_0}).$$

Here and in what follows, *ffun* is the (filtered) A_∞-category of filtered functors, and $\hom^{\leq 0}_{\mathit{ffun}}$ stands for the morphisms in that category that do not shift filtration, namely

the natural transformations (between filtered functors) that preserve filtrations. Furthermore, all the cohomological identities from [53, Section (10a), p. 134] between these natural transformations should now hold in the 0-categories

$$(H^0(\mathit{ffun}(\mathcal{A}^{i'}, \mathcal{A}^{i''})))_0 \tag{3.16}$$

(i.e. persistence level 0) of the persistence homological categories $H^0(\mathit{ffun}(\mathcal{A}^{i'}, \mathcal{A}^{i''}))$. (In (3.16), the 0-superscript means cohomological degree 0 and the 0-subscript stands for the 0 persistence level subcategory.)

We will sometimes refer to $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$ as a family of A_∞-categories over $\mathcal{I}$ and call $\mathcal{I}$ the base of the family. Similarly, in case we have a coherent system on $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$, we will call it a coherent system over $\mathcal{I}$.

One way to assemble a coherent system out of a family of A_∞-categories $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$ is first to try to embed all of them into one total A_∞-category $\mathcal{A}^i \subset \mathcal{A}^{\mathrm{tot}}$ by quasi-equivalences and then seek for suitable projection functors $\mathcal{A}^{\mathrm{tot}} \to \mathcal{A}^j$ for all i. The functors $\mathcal{F}^{i_1,i_0} : \mathcal{A}^{i_0} \to \mathcal{A}^{i_1}$ participating in the coherent system will then be defined as the composition of the inclusions $\mathcal{A}^{i_0} \subset \mathcal{A}^{\mathrm{tot}}$ with the projections $\mathcal{A}^{\mathrm{tot}} \to \mathcal{A}^{i_1}$. We will soon adapt this scheme to the filtered framework. But first we need to introduce some relevant notions.

We begin with A_n-categories and functors. An A_n-category $\mathcal{A}$ is the same as an A_∞-category with the exception that now we have $\mu_k^{\mathcal{A}}$ only for $k = 1, \dots, n$. The $\mu_k^{\mathcal{A}}$'s are required to satisfy the subset of the A_∞-identities that involve only the $\mu_k^{\mathcal{A}}$'s with $1 \le k \le n$. In case the category $\mathcal{A}$ is clear from the context we will sometimes omit the superscript from $\mu_k^{\mathcal{A}}$.

Let $\mathcal{A}$ and $\mathcal{B}$ be two A_n-categories, and let $m \le n$. Similarly to A_∞-functors, we have A_m-functors $\mathcal{F} : \mathcal{A} \to \mathcal{B}$. They are defined in the same way as A_∞-functors, but now the higher-order components $\mathcal{F}_k$ of $\mathcal{F}$ are defined only for $1 \le k \le m$. The $\mathcal{F}_k$'s, $\mu_i^{\mathcal{A}}$'s, and $\mu_j^{\mathcal{B}}$'s are required to satisfy the same identities as for A_∞-functors, involving only indices with $1 \le k \le m$ and $i, j \le n$. A_n-functors can be composed in a similar way as A_∞-functors. Finally, the notions of pre-natural and natural transformations between A_∞-functors generalize to A_m-functors in a similar way.

Let $\mathcal{A}$ be an A_n-category, with $n \ge 3$. Let $X \in \mathrm{Obj}(\mathcal{A})$ and $2 \le k \le n$. An element $e_X \in \hom_{\mathcal{A}}(X, X)$ is called a strict A_k-unit if it satisfies the following conditions: e_X is a cycle; for all $X' \in \mathrm{Obj}(\mathcal{A})$, the maps

$$\mu_2(-, e_X) : \hom_{\mathcal{A}}(X', X) \to \hom_{\mathcal{A}}(X', X),$$
$$\mu_2(e_X, -) : \hom_{\mathcal{A}}(X, X') \to \hom_{\mathcal{A}}(X, X')$$

are the identity maps; and moreover $\mu_j(-, \dots, -, e_X, -, \dots, -) = 0$ for all $2 < j \le k$. An A_n-category is called strictly A_k-unital if we are given (as part of the structure) strict A_k-units e_X for every object $X \in \mathrm{Obj}(\mathcal{A})$. If an A_n-category is strictly A_n-unital (i.e., $k = n$) we will simply say that it is strictly unital.

Let $\mathcal{A}$, $\mathcal{B}$ be two A_n-categories which are both strictly A_k-unital. An A_m-functor $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ is called strictly A_l-unital for $l \leq m$ if $\mathcal{F}_1(e_X) = e_{\mathcal{F}(X)}$ for all $X \in \mathrm{Obj}(\mathcal{A})$ and $\mathcal{F}_i(-, \ldots, -, e_X, -, \ldots, -) = 0$ for every $2 \leq i \leq l$.

Similarly to strict units, we also have the notion of strict isomorphisms. Let $\mathcal{A}$ be an A_n-category ($3 \leq n$) which is strictly A_l-unital ($l \leq n$), and denote its strict units by $e_Z \in \hom_{\mathcal{A}}(Z, Z)$, $Z \in \mathrm{Obj}(\mathcal{A})$. Let $X, Y \in \mathrm{Obj}(\mathcal{A})$ and let $u \in \hom_{\mathcal{A}}(X, Y)$ be a cycle. We say that u is a strict A_k-isomorphism (where $k \leq n$) if there exists a cycle $v \in \hom_{\mathcal{A}}(Y, X)$ such that

$$\mu_2(u, v) = e_X, \quad \mu_2(v, u) = e_Y, \quad \mu_j(-, \ldots, -, u_X, -, \ldots, -) = 0 \ \forall\, 2 < j \leq k.$$

If $k = n$, we will simply say that u is a strict isomorphism.

Remark 3.15. (a) In what follows, we will view A_∞-categories as a special case of A_n-categories by allowing $n = \infty$ (of course, one needs to slightly adjust the definition, since for $n = \infty$ the operations μ_k exist only for $1 \leq k < n$ and not for $1 \leq k \leq n$). The same remark applies also to A_n-functors, (pre-)natural transformations, and strict units, and we will view their A_∞-counterparts as a special case of the respective A_n-objects.

(b) A_n-categories are also $A_{n'}$-categories for all $1 \leq n' \leq n$; therefore, we will sometimes reduce n to the minimal value that is relevant in the context. A similar remark applies also to functors, (pre-)natural transformations, and strict units.

(c) Let $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ be an A_n-functor and $n' \leq n$. We denote by $\{\mathcal{F}\}_{n'}$ the $A_{n'}$-functor obtained from $\mathcal{F}$ by ignoring the terms of order greater than n' (i.e., the terms $\mathcal{F}_k$ for $n' < k \leq n$). We call $\{\mathcal{F}\}_{n'}$ the $A_{n'}$-reduction of $\mathcal{F}$.

In what follows, we will be mainly interested in A_2-functors $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ between two A_3-categories. Unwrapping the above definitions, in this case this means that $\mathcal{F}_1 : \hom_{\mathcal{A}}(X_0, X_1) \to \hom_{\mathcal{B}}(\mathcal{F}X_0, \mathcal{F}X_1)$ is a chain map and that $\mathcal{F}_1$ respects composition of morphisms up to a chain homotopy given by $\mathcal{F}_2$. In other words, $\mathcal{F}_2 : \hom_{\mathcal{A}}(X_0, X_1) \otimes \hom_{\mathcal{A}}(X_1, X_2) \to \hom_{\mathcal{A}}(X_0, X_2)$ defines a chain homotopy between $\mu_2 \circ (\mathcal{F}_1 \otimes \mathcal{F}_1)$ and $\mathcal{F}_1 \circ \mu_2$. Passing to homology, $\mathcal{F}$ induces a (non-unital) functor $H(\mathcal{F}) : H(\mathcal{A}) \to H(\mathcal{B})$ between the homological categories $H(\mathcal{A})$ and $H(\mathcal{B})$. Note that $H(\mathcal{A})$ and $H(\mathcal{B})$ are genuine categories, since $\mathcal{A}$ and $\mathcal{B}$ were assumed to be A_3-categories.

Assuming $\mathcal{A}$ and $\mathcal{B}$ to be A_2-unital, some of our A_2-functors $\mathcal{F}$ will be strictly A_1-unital, which means that $\mathcal{F}_1(e_X) = e_{\mathcal{F}X}$ for all objects X. In particular, this implies that $\mathcal{F}$ is homologically unital (i.e., the functor $H(\mathcal{F})$ is unital).

We now turn to the filtered setting. Filtered A_n-categories, A_m-functors and their pre-natural transformations are defined precisely in the same way as their filtered A_∞-counterparts. Strict A_k-units are required by definition to be in filtration level 0, and the same goes for strict isomorphisms. (Below we will not attach the adjective

"filtered" to units/isomorphisms, implicitly assuming that whenever we are in the filtered setting these elements are in filtration level 0.)

If $\mathcal{A}$ is a filtered A_3-category which is strictly A_2-unital then $H(\mathcal{A})$ is a persistence category. If $\mathcal{F}$ is a filtered A_2-functor which is A_1-unital then $H(\mathcal{F})$ is a persistence functor.

Finally, we need the notion of a filtered quasi-equivalence. Let $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ be an A_2-functor between filtered A_3-categories. We call $\mathcal{F}$ a filtered quasi-equivalence if $\mathcal{F}$ is a filtered functor and its homological functor $H(\mathcal{F}) : H(\mathcal{A}) \to H(\mathcal{B})$ is an equivalence of persistence categories.

Next, we need to introduce persistence Hochschild cohomology. Let $\mathcal{A}$, $\mathcal{B}$ be A_3-categories which are strictly A_2-unital. Let $\mathcal{F}, \mathcal{G}$ be filtered A_2-functors which are A_1-unital. Given $X_0, \dots, X_r \in \mathrm{Obj}(\mathcal{A})$, we will abbreviate $\mathcal{A}(X_0, \dots, X_r) := \hom_{\mathcal{A}}(X_0, X_1) \otimes \cdots \otimes \hom_{\mathcal{A}}(X_{r-1}, X_r)$ and view it as a filtered chain complex in the standard way (and similarly for $\mathcal{B}$). The persistence Hochschild cochain complex associated to the above data is a cochain complex of persistence modules. The α-persistence level of this cochain complex in degree r is

$$\mathrm{PCC}^{r;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G}) := \prod_{X_0,\dots,X_r} H^*\Big(\hom_{\mathbf{k}}^{\leq\alpha}\big(\mathcal{A}(X_0, \dots, X_r), \mathcal{B}(\mathcal{F}X_0, \mathcal{G}X_r)\big)\Big), \tag{3.17}$$

where the product runs over all r-tuples of objects in $\mathcal{A}$. Here we view

$$\hom_{\mathbf{k}}\big(\mathcal{A}(X_0, \dots, X_r), \mathcal{B}(\mathcal{F}X_0, \mathcal{G}X_r)\big)$$

as a filtered cochain complex, endowed with the standard differential, which will be denoted below by δ. The α-level filtration on this cochain complex, denoted here by $\hom_{\mathbf{k}}^{\leq\alpha}$, consists of those (graded) homomorphisms that shift filtration by at most α. Finally, H^* stands for persistence cohomology (in all degrees). Note that the spaces PCC are in fact bigraded since the $\hom_{\mathbf{k}}$ term is graded in itself. Thus a more detailed description of (3.17) would be to define $\mathrm{PCC}^{r,s}$ with the term H^* replaced by cohomology H^s in degree s. However, whenever it is not necessary we will ignore the s-degree.

The differential

$$\partial_{\mathrm{PCC}} : \mathrm{PCC}^{r;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G}) \to \mathrm{PCC}^{r+1;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G})$$

is defined as follows. Let $T \in \hom_{\mathbf{k}}^{\leq\alpha}(\mathcal{A}(X_0, \dots, X_r), \mathcal{B}(\mathcal{F}X_0, \mathcal{G}X_r))$ be a δ-cocycle (i.e., a chain map). Then $\partial_{\mathrm{CC}}([T])$ is defined as the δ-cohomology class $[S]$ of the

chain map

$$\begin{aligned} S(a_1,\dots,a_{r+1}) := {} & \epsilon'\mu_2(\mathcal{F}_1(a_1), T(a_2,\dots,a_{r+1})) \\ & + \epsilon''\mu_2(T(a_1,\dots,a_r), \mathcal{G}_1(a_{r+1})) \\ & + \epsilon \sum_{k=0}^{r} T(a_1,\dots,\mu_2(a_{k+1},a_{k+2}),\dots,a_{r+1}), \end{aligned} \tag{3.18}$$

where $\epsilon', \epsilon'', \epsilon = \pm 1$ are signs that depend on the degrees of T, the a_i's, and r; see [53, Chapter I, Section (1f)]. Since we work over $\mathbf{k} = \mathbb{Z}_2$, these will play no role in our considerations anyway.

The cohomology of $(\mathrm{PCC}, \partial_{\mathrm{PCC}})$ is a graded persistence module, which we call the persistence Hochschild cohomology of $(\mathcal{A}, \mathcal{B})$ with respect to the functors $\mathcal{F}, \mathcal{G}$. We denote it by $\mathrm{PHH}^{*;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G})$, $\alpha \in \mathbb{R}$.

Apart from PHH, there is a persistence variant of the classical Hochschild cohomology for the categories $H(\mathcal{A})$ and $H(\mathcal{B})$ with respect to the homological functors $[\mathcal{F}]$, $[\mathcal{G}]$. This is defined in the same way as in [53, Chapter I, Section (1f)], only that we also take the persistence structure into account. We will use the following notation. We write

$$H^{\mathcal{A}}(X, Y) := H(\hom_{\mathcal{A}}(X, Y))$$

for the persistence homology (in all degrees) of $\hom_{\mathcal{A}}(X, Y)$ with respect to $\mu_1^{\mathcal{A}}$, and similarly for $\mathcal{B}$. Set also

$$H^{\mathcal{A}}(X_0,\dots,X_r) := H^{\mathcal{A}}(X_0, X_1) \otimes \cdots \otimes H^{\mathcal{A}}(X_{r-1}, X_r),$$

viewed as a (graded) persistence module. Set $A := H(\mathcal{A})$ and $B := H(\mathcal{B})$, the homological categories of $\mathcal{A}$ and $\mathcal{B}$, and set $F := [\mathcal{F}]$ and $G := [\mathcal{G}]$ to be the homological functors corresponding to $\mathcal{F}$ and $\mathcal{G}$. The cochain complex $\overline{\mathrm{PCC}}$ for the Hochschild cohomology of $H(\mathcal{A})$ and $H(\mathcal{B})$ has the following persistence module structure in degree r:

$$\overline{\mathrm{PCC}}^{r;\leq\alpha}(A, B; F, G) := \prod_{X_0,\dots,X_r} \Big(\hom_{\mathrm{per}}^{\leq\alpha}\big(H^{\mathcal{A}}(X_0,\dots,X_r), H^{\mathcal{B}}(\mathcal{F}X_0, \mathcal{G}X_r)\big)\Big), \tag{3.19}$$

where α denotes the persistence level and $\hom_{\mathrm{per}}^{\leq\alpha}$ stands for (graded) homomorphisms from the persistence module $H^{\mathcal{A}}(X_0,\dots,X_r)$ to the persistence module $H^{\mathcal{B}}(\mathcal{F}X_0, \mathcal{G}X_r)[\alpha]$. We are using the notation and conventions from Section 2.2.4 (see (2.15)). Namely, $H^{\mathcal{B}}(\mathcal{F}X_0, \mathcal{G}X_r)[\alpha]$ stands for $H^{\mathcal{B}}(\mathcal{F}X_0, \mathcal{G}X_r)$ shifted by α in terms of the persistence parameter. The structure map from level $\hom_{\mathrm{per}}^{\leq\alpha}$ to level $\hom_{\mathrm{per}}^{\leq\beta}$, whenever $\alpha \leq \beta$, is given by composing with the structure map indexed by α, β in the persistence module $H^{\mathcal{B}}(\mathcal{F}X_0, \mathcal{G}X_r)$. The differential $\partial_{\overline{\mathrm{CC}}}$ has a similar

expression to (3.18), and an explicit formula can be found in [53, Chapter I, Section (1f), p. 13]. We denote the resulting persistence cohomology by $\overline{\mathrm{PHH}}^{*;\leq\alpha}(A, B; F, G)$, $\alpha \in \mathbb{R}$.

Remark 3.16. (a) It is easy to see that the data required to define the second version, $\overline{\mathrm{PHH}}$, of Hochschild cohomology (and in fact even $\overline{\mathrm{PCC}}$) is entirely homological. In fact, this cohomology can be defined for every pair of persistence categories and a pair of persistence functors between them. In contrast, PCC and PHH seem to depend on some chain level information from $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{F}, \mathcal{G}$.

The two cochain complexes PCC and $\overline{\mathrm{PCC}}$ are related to each other by a Künneth-type short exact sequence of persistence modules:

$$0 \longrightarrow E_{r-1}(H(\mathcal{A}), H(\mathcal{B})) \longrightarrow \mathrm{PCC}^r(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G}) \longrightarrow \overline{\mathrm{PCC}}^r(A, B; F, G) \longrightarrow 0, \tag{3.20}$$

where $E_{r-1}(-,-)$ is a derived functor of $\hom_{\mathrm{per}}$ (which is analogous to Ext^1) in the category of persistence modules. More specifically, when inserting the components of $X_0, \ldots, X_r$ from (3.17) and (3.19) into (3.20), the term E_{r-1} involves only homologies of the type $H_*(\mathcal{A}(X_0, \ldots, X_r))$ and $H_*^{\mathcal{B}}(\mathcal{F}X_0, \mathcal{G}X_r)$ of total degree $r-1$. The former term $H_*(\mathcal{A}(X_0, \ldots, X_r))$ can also be related, via short exact sequences (a persistence analog of the universal coefficients theorem), to $H_*^{\mathcal{A}}(X_0, X_1) \otimes \cdots \otimes H_*^{\mathcal{A}}(X_{r-1}, X_r)$ and Tor_1-like derived functors associated to tensor products of persistence modules. We refer the reader to [13, Section 8] for the precise formulation of the Künneth and universal coefficients theorems for persistence homology (see also [49]).

(b) Let $3 \leq n \leq \infty$, $m \leq n$. Given two A_n-categories $\mathcal{A}$, $\mathcal{B}$ and two A_m-functors $\mathcal{F}, \mathcal{G} : \mathcal{A} \to \mathcal{B}$, the filtered natural transformations $\hom_{ffun(\mathcal{A},\mathcal{B})}(\mathcal{F}, \mathcal{G})$ form a filtered chain complex (by filtered natural transformations we mean those that respect filtrations up to a bounded shift). The level-α filtration $\hom^{\leq\alpha}_{ffun(\mathcal{A},\mathcal{B})}$ consists of those natural transformations that shift filtration by at most α. Apart from this filtration, this space admits yet another filtration called the length filtration. This one is indexed by the natural numbers and is decreasing. Its p-level $F^p \hom^{\leq\alpha}_{ffun(\mathcal{A},\mathcal{B})}(\mathcal{F}, \mathcal{G})$ is given by those natural transformations $T \in \hom^{\leq\alpha}_{ffun(\mathcal{A},\mathcal{B})}(\mathcal{F}, \mathcal{G})$ with $T_0 = \cdots = T_p = 0$. There is a spectral sequence of persistence modules associated to this filtration. A simple calculation shows that its first page, at persistence level α, is given by

$$E_1^{r,s;\leq\alpha} = \mathrm{PCC}^{r,s;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G}).$$

Here we have used the second degree s on PCC as briefly explained earlier. See [53, Chapter I, Section (1f), p. 13] for more details in the non-filtered case. The filtration F^p is bounded in the case of A_m-functors with m finite. For A_∞-functors it is not bounded, and its associated spectral sequence might not converge; however, following [53] we will use this sequence as a tool for comparing the homologies of different $\hom_{ffun}$'s.

Let $\mathcal{A}, \mathcal{B}, \mathcal{C}$ be filtered A_n-categories ($3 \leq n \leq \infty$). Let $2 \leq m \leq n$, and denote by *nu-ffun*$(\mathcal{C}, \mathcal{A})$ and *nu-ffun*$(\mathcal{C}, \mathcal{B})$ the filtered A_m-categories of non-unital (or, better said, not necessarily unital) filtered A_m-functors $\mathcal{C} \to \mathcal{A}$ and $\mathcal{C} \to \mathcal{B}$, respectively. Let $\mathcal{G} : \mathcal{A} \to \mathcal{B}$ be a non-unital filtered A_m-functor and denote by

$$\mathcal{L}_{\mathcal{G}} : \textit{nu-ffun}(\mathcal{C}, \mathcal{A}) \to \textit{nu-ffun}(\mathcal{C}, \mathcal{B})$$

the functor induced by left composition with $\mathcal{G}$. Note that this is a (non-unital) filtered A_m-functor. An immediate consequence of Remark 3.16 is the following.

Lemma 3.17 (cf. [53, Lemma 1.7]). *If $\mathcal{G}$ is homologically full and faithful in the persistence sense, then so is $\mathcal{L}_{\mathcal{G}}$.*

The next lemma deals with invariance of persistence Hochschild cohomology under filtered quasi-equivalences. It is a persistence analog of a very special case of [53, Lemma 2.6], with several additional very strong assumptions made in order to accommodate the persistence case.

Lemma 3.18. *Let $\mathcal{A}$ be a filtered A_n-category ($3 \leq n \leq \infty$) with strict A_3-units, and let $\widetilde{\mathcal{A}} \subset \mathcal{A}$ be a full subcategory such that the inclusion $\mathcal{J} : \widetilde{\mathcal{A}} \to \mathcal{A}$ is a filtered quasi-equivalence. Suppose that $\mathcal{P} : \mathcal{A} \to \widetilde{\mathcal{A}}$ is a filtered A_2-functor which is A_1-unital, and assume that $\mathcal{P} \circ \mathcal{J} = \mathbb{1}_{\widetilde{\mathcal{A}}}$ (as A_2-functors). Assume further that for every $X \in \mathrm{Obj}(\mathcal{A})$ we have a strict isomorphism $u_X \in \hom_{\mathcal{A}}(X, \mathcal{P}(X))$. Then the map induced by the restriction*

$$\rho : \mathrm{PHH}(\mathcal{A}, \widetilde{\mathcal{A}}; \mathcal{P}, \mathcal{P}) \to \mathrm{PHH}(\widetilde{\mathcal{A}}, \widetilde{\mathcal{A}}; \mathbb{1}_{\widetilde{\mathcal{A}}}, \mathbb{1}_{\widetilde{\mathcal{A}}})$$

is a (bigraded) isomorphism of persistence modules.

We omit the proof since it is very similar to that of [53, Lemma 2.6]. The role of the strict isomorphisms u_X is to facilitate the definition, in the framework of PCC and PHH, of certain chain maps and chain homotopies that appear in the original proof. Note also that [53, Lemma 2.6] is stated for any two functors; here we only need it for the functor $\mathcal{P}$, which simplifies things further.

We now get to extending filtered A_n-functors from a subcategory to a larger one. The following lemma is a persistence analog, this time of a special case of [53, Lemma 1.10], and again with several additional assumptions.

Lemma 3.19. *Let $\mathcal{A}$, $\mathcal{B}$ be filtered A_n-categories ($3 \leq n \leq \infty$) and $\widetilde{\mathcal{A}} \subset \mathcal{A}$ a full subcategory. Let $\mathcal{P} : \mathcal{A} \to \mathcal{B}$ be a filtered A_2-functor and denote by $\widetilde{\mathcal{P}} := \mathcal{P}|_{\widetilde{\mathcal{A}}}$ its restriction to $\widetilde{\mathcal{A}}$. Assume that the map*

$$\rho : \mathrm{PHH}^r(\mathcal{A}, \mathcal{B}; \mathcal{P}, \mathcal{P}) \to \mathrm{PHH}^r(\widetilde{\mathcal{A}}, \mathcal{B}; \widetilde{\mathcal{P}}, \widetilde{\mathcal{P}})$$

induced by the restriction is an isomorphism of persistence modules for every r. Then every filtered A_n-functor $\widetilde{\mathcal{Q}} : \widetilde{\mathcal{A}} \to \mathcal{B}$ with $\{\widetilde{\mathcal{Q}}\}_2 = \widetilde{\mathcal{P}}$ can be extended to a filtered A_n-functor $\mathcal{Q} : \mathcal{A} \to \mathcal{B}$ with $\{\mathcal{Q}\}_2 = \mathcal{P}$. (See point (c) *of Remark 3.15.)*

We omit the proof again since it is very similar to that of [53, Lemma 1.10], with straightforward modifications needed for the persistence setting.

We are now ready to assemble a coherent system of filtered A_∞-categories out of a family of categories that are all included into one total category. Let $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$ be a family of filtered, strictly unital A_∞-categories over $\mathcal{I}$. Suppose there is a filtered strictly unital A_∞-category $\mathcal{A}^{\mathrm{tot}}$ such that for every $i \in \mathcal{I}$, $\mathcal{A}^i$ is a full subcategory of $\mathcal{A}^{\mathrm{tot}}$. Denote by $\mathcal{J}^i : \mathcal{A}^i \to \mathcal{A}^{\mathrm{tot}}$ the inclusion functor and assume that $\mathcal{J}^i$ is a filtered quasi-equivalence. We will refer to a category $\mathcal{A}^{\mathrm{tot}}$ as above, together with the inclusion functors $\mathcal{J}^i$, as a (filtered, strictly unital) *total A_∞-category over $\mathcal{I}$*.

Assume that for every $i \in \mathcal{I}$ there is a filtered A_2-functor $\mathcal{P}r^i : \mathcal{A}^{\mathrm{tot}} \to \mathcal{A}^i$ which is strictly A_1-unital and such that $\mathcal{P}r^i \circ \mathcal{J}^i = \mathbb{1}_{\mathcal{A}^i}$ as A_2-functors. Assume further that the following holds for every $i \in \mathcal{I}$: for every $X \in \mathrm{Obj}(\mathcal{A}^{\mathrm{tot}})$ there exists a strict isomorphism $u^i_X \in \hom_{\mathcal{A}^{\mathrm{tot}}}(X, \mathcal{P}r^i(X))$.

Proposition 3.20. *Under the above assumptions each of the A_2-functors $\mathcal{P}r^i$, $i \in \mathcal{I}$, can be extended to a filtered A_1-unital A_∞-functor $\mathcal{Q}^i : \mathcal{A}^{\mathrm{tot}} \to \mathcal{A}^i$ (i.e., $\{\mathcal{Q}^i\}_2 = \mathcal{P}r^i$). Moreover, the functors $\mathcal{F}^{j,i} := \mathcal{Q}^j \circ \mathcal{J}^i : \mathcal{A}^i \to \mathcal{A}^j$, $i, j \in \mathcal{I}$, are filtered, A_1-unital, and form a coherent system of filtered A_∞-categories over $\mathcal{I}$. The filtered natural transformations $T^{i_2,i_1,i_0} : \mathcal{F}^{i_2,i_1} \circ \mathcal{F}^{i_1,i_0} \to \mathcal{F}^{i_2,i_0}$ will be described in the proof.*

Proof. The restriction of the A_2-functors $\mathcal{P}r^i$ to the subcategory $\mathcal{A}^i \subset \mathcal{A}^{\mathrm{tot}}$ is the A_2-reduction of the identity A_∞-functor on $\mathcal{A}^i$. Using Lemmas 3.18 and 3.19 we can extend $\mathcal{P}r^i$ to the desired A_∞-functor $\mathcal{Q}^i$.

The construction of the natural transformations $T^{i_2,i_1,i_0} : \mathcal{F}^{i_2,i_1} \circ \mathcal{F}^{i_1,i_0} \to \mathcal{F}^{i_2,i_0}$ follows the same scheme as in [53, Chapter II, Section (10a), p. 134].

Consider the homological functor induced by left composition with $\mathcal{Q}^i$, viewed as a persistence functor:

$$H(\mathcal{L}_{\mathcal{Q}^i}) : H(\mathit{ffun}(\mathcal{A}^{\mathrm{tot}}, \mathcal{A}^{\mathrm{tot}})) \to H(\mathit{ffun}(\mathcal{A}^{\mathrm{tot}}, \mathcal{A}^i)).$$

Recall that $\mathcal{Q}^i$ is a filtered quasi-equivalence and, in particular, homologically full and faithful in the persistence sense. By Lemma 3.17, the action of the functor $H(\mathcal{L}_{\mathcal{Q}^i})$ on morphisms is an isomorphism of persistence modules:

$$H(\mathcal{L}_{\mathcal{Q}^i}) : \hom_{H(\mathit{ffun}(\mathcal{A}^{\mathrm{tot}},\mathcal{A}^{\mathrm{tot}}))}(\mathcal{J}^i \circ \mathcal{Q}^i, \mathbb{1}_{\mathcal{A}^{\mathrm{tot}}}) \to \hom_{H(\mathit{ffun}(\mathcal{A}^{\mathrm{tot}},\mathcal{A}^i))}(\mathcal{Q}^i, \mathcal{Q}^i).$$

Let $S^i \in \hom^{\leq 0}_{\mathit{ffun}(\mathcal{A}^{\mathrm{tot}},\mathcal{A}^{\mathrm{tot}})}(\mathcal{J}^i \circ \mathcal{Q}^i, \mathbb{1}_{\mathcal{A}^{\mathrm{tot}}})$ be a cycle whose homology class $[S^i]$ is sent by $H(\mathcal{L}_{\mathcal{Q}^i})$ to $[\mathbb{1}] \in H(\hom_{\mathit{ffun}(\mathcal{A}^{\mathrm{tot}},\mathcal{A}^i)}(\mathcal{Q}^i, \mathcal{Q}^i))$. Having defined S^i as above for all $i \in \mathcal{I}$, we define

$$T^{i_2,i_1,i_0} := \mathcal{L}_{\mathcal{Q}^{i_2}}(\mathcal{R}_{\mathcal{J}^{i_0}}(S^{i_1})),$$

where $\mathcal{R}_{\mathcal{J}^{i_0}}$ is the functor $\mathit{ffun}(\mathcal{A}^{\mathrm{tot}}, \mathcal{A}^{\mathrm{tot}}) \to \mathit{ffun}(\mathcal{A}^{i_0}, \mathcal{A}^{\mathrm{tot}})$ induced by right composition with $\mathcal{J}^{i_0}$. ∎

3.3.7 Proof of Theorem 3.12, part 3: Coherent systems of Fukaya categories

In order to apply the algebraic statements from Section 3.3.6, in particular Proposition 3.20, to the case of Fukaya categories, we need two additional ingredients coming from geometry. First, we need to construct a filtered total Fukaya category $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X})$ that contains all the Fukaya categories $\mathcal{F}uk(\mathcal{X}; i)$, constructed via various perturbation data i, as quasi-equivalent subcategories of $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X})$. Second, we need to construct filtered A_2-functors $\mathcal{P}r^i : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}) \to \mathcal{F}uk(\mathcal{X}; i)$ that are left inverses of the inclusions $\mathcal{F}uk(\mathcal{X}; i) \subset \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X})$.

We begin with the first ingredient, namely the construction of $\mathcal{F}uk^{\mathrm{tot}}$. To this end, recall the spaces of perturbation data $\mathcal{B}(\bar{\mathcal{X}}; R, r)$ from page 132. We fix the symplectic embedding ϕ and the parameters $R_{\bar{\mathcal{X}}}$ and $r \leq R \leq R_{\bar{\mathcal{X}}}$. We will outline now a construction of a filtered strictly unital total A_∞-category over $\mathcal{B}(\bar{\mathcal{X}}; R, r)$. (See page 140 for the meaning of a total category over a base.) More specifically, we will construct a filtered strictly unital A_∞-category $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$ with the following property. For every $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$, the filtered Fukaya category $\mathcal{F}uk(\mathcal{X}; i)$ is a full subcategory of $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$ and the inclusion $\mathcal{J}^i : \mathcal{F}uk(\mathcal{X}; i) \to \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$ is a filtered quasi-equivalence.

The construction of $\mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}}; R, r)$ follows steps similar to those of the construction introduced in [53, Chapter II, Section (10a), pp. 134–135], with some significant modifications necessary to accommodate the filtered setting. The objects of $\mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}}; R, r)$ are pairs (L, i), where $L \in \mathcal{X}$ and $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$ is a choice of admissible perturbation data. The morphism space between (L_0, i_0) and (L_1, i_1) is defined to be the Floer complex $\mathrm{CF}(L_0, L_1; \mathcal{D}_{(L_0,i_0),(L_1,i_1)})$, where the Floer datum $\mathcal{D}$ is defined as in our earlier construction of filtered Fukaya categories, with the only restriction that if $i_0 = i_1$ then the Floer datum $\mathcal{D}_{(L_0,i_0),(L_1,i_1)}$ should agree with that of (L_0, L_1) in the category $\mathcal{F}uk(\mathcal{X}; i)$. Another important point is that we require the Floer data for pairs of the type $((L, i), (L, j))$ to continue to be of the same type as in our construction of filtered Fukaya categories. Namely, we take here a pair of a Morse function and a Riemannian metric, such that the Morse function has a unique local maximum (i.e., a unique critical point of index $n = \dim_{\mathbb{C}} X$).

The next step in the construction of the total category is to choose consistent perturbation data $\mathcal{P}^{\mathrm{tot}} = \mathcal{P}^{\mathrm{tot}}(\mathcal{X}; R, r)$ with the restriction analogous to the one imposed on the Floer data. Namely, whenever we have a cluster of punctured disks decorated entirely by pairs (L_k, i) with the same i, the value of $\mathcal{P}^{\mathrm{tot}}$ on such a cluster coincides with the one prescribed by the perturbation data i. Apart from that, there will be one important difference from the way we defined the perturbation data for each $\mathcal{F}uk(\mathcal{X}; i)$, which we now describe. Given a tuple $(i_0, \ldots, i_d)$ with $i_k \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$,

we define the following two quantities:

$$\begin{aligned} R^{(i_0,\dots,i_d)} &:= \sup\{\widetilde{R} \mid \widetilde{R} \leq R_{\overline{\mathcal{X}}} \text{ and } i_k \in \mathcal{B}(\overline{\mathcal{X}}; \widetilde{R}, \widetilde{R}) \ \forall\, 0 \leq k \leq d\}, \\ r^{(i_0,\dots,i_d)} &:= \inf\{\widetilde{r} \mid 0 \leq \widetilde{r} \text{ and } i_k \in \mathcal{B}(\overline{\mathcal{X}}; \widetilde{r}, \widetilde{r}) \ \forall\, 0 \leq k \leq d\}. \end{aligned} \tag{3.21}$$

In other words, $R^{(i_0,\dots,i_d)}$ measures the supremal radius of the sub-balls in the embedding ϕ on which the almost complex structures in all the perturbation data i_k are standard. The other quantity, $r^{(i_0,\dots,i_d)}$, measures the infimal upper bound on the perturbation 1-forms in all the perturbation data i_k. Note that since $i_k \in \mathcal{B}(\overline{\mathcal{X}}; R, r)$ for all k, we have

$$r^{(i_0,\dots,i_d)} \leq r \leq R \leq R^{(i_0,\dots,i_d)}. \tag{3.22}$$

Turning to the definition of the perturbation data $\mathcal{P}^{\mathrm{tot}}$, we require it to satisfy the following conditions. Let S_{d+1} be a $(d+1)$-punctured disk, decorated by the tuple $((L_0, i_0), \dots, (L_d, i_d))$, and denote by (K, J) the value of the perturbation data on S_{d+1}. Recall the symplectic embedding ϕ from (3.10) and the balls B_j (see page 129). We require that:

(1) $K|_{\phi(B_j(R^{(i_0,\dots,i_d)}))} \equiv 0$ for every j. Here $B_j(R^{(i_0,\dots,i_d)}) \subset B_j$ denotes the smaller ball of radius $R^{(i_0,\dots,i_d)}$.

(2) $J_z|_{\phi(B_j(R^{(i_0,\dots,i_d)}))} = \phi_*(J_{\mathrm{std}})$ for every j and every $z \in S_{d+1}$.

In addition to the above two conditions, we also require that

$$\nu(\mathcal{P}^{\mathrm{tot}}((L_0, i_0), \dots, (L_d, i_d))) < C(d+1)(r^{(i_0,\dots,i_d)})^2. \tag{3.23}$$

Here C is the constant from (3.11), and, similarly to (3.8),

$$\nu(\mathcal{P}^{\mathrm{tot}}((L_0, i_0), \dots, (L_d, i_d))) := C_{d+1} \sup \|K(\mathcal{P}^{\mathrm{tot}}, S_{d+1})\|,$$

where now the supremum goes over all $(d+1)$-punctured disks S_{d+1} that are decorated by $((L_0, i_0), \dots, (L_d, i_d))$. Finally, the above requirements extend to clusters of punctured disks in a similar way.

We claim that there exists a consistent choice of perturbation data $\mathcal{P}^{\mathrm{tot}}$ satisfying the above conditions. The proof is similar to that of the corresponding statement for the perturbation data $i \in \mathcal{B}(\overline{\mathcal{X}}; R, r)$.

With a choice of perturbation data as above, one can define an A_∞-category in the same way as we defined our earlier Fukaya categories. We denote this category by $\mathcal{F}uk^{\mathrm{tot}}(\overline{\mathcal{X}}; R, r)$, or sometimes by $\mathcal{F}uk^{\mathrm{tot}}(\overline{\mathcal{X}}; \mathcal{P}^{\mathrm{tot}})$ when we want to emphasize the choice of the perturbation data $\mathcal{P}^{\mathrm{tot}}$ used to define it.

We claim that $\mathcal{F}uk^{\mathrm{tot}}(\overline{\mathcal{X}}; R, r)$ is a filtered and strictly unital A_∞-category. This follows by the same arguments we used for $\mathcal{F}uk(\mathcal{X}; i)$, $i \in \mathcal{B}(\overline{\mathcal{X}}; R, r)$, with minor modifications. The important points are that the analogs of inequalities (3.11), (3.12), and (3.13) will continue to hold with R and r replaced by $R^{(i_0,\dots,i_d)}$ and $r^{(i_0,\dots,i_d)}$

respectively, and $\nu(\mathcal{P}_{d+1})$ by $\nu(\mathcal{P}^{\mathrm{tot}}((L_0,i_0),\dots,(L_d,i_d)))$. This completes the outline of the construction of the category $\mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}};R,r)$.

For every $i \in \mathcal{B}(\bar{\mathcal{X}};R,r)$ there is an obvious inclusion

$$\mathcal{J}^i : \mathcal{F}uk(\mathcal{X};i) \to \mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}};R,r).$$

Clearly, this functor is filtered, and we claim that it is a filtered quasi-equivalence. To see the latter, first note that by construction $\mathcal{J}^i$ is full and faithful. Now, any object $(L,k) \in \mathrm{Obj}(\mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}};R,r))$ is isomorphic to (L,i) via an isomorphism that lies in $\hom^{\leq 0}$ of $\mathcal{F}uk^{\mathrm{tot}}$. This follows from Morse theory, since the hom between (L,k) and (L,i) is the Morse complex of L with respect to a Morse function with a unique local maximum. This shows that $\mathcal{J}^i$ is a filtered quasi-equivalence.

Next, for every $i \in \mathcal{B}(\bar{\mathcal{X}};R,r)$, we construct a filtered A_2-functor

$$\mathcal{P}r^i : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r) \to \mathcal{F}uk(\mathcal{X};i),$$

which is strictly A_1-unital and satisfies $\mathcal{P}r^i \circ \mathcal{J}^i = \mathbb{1}_{\mathcal{F}uk(\mathcal{X};i)}$.

Let $i \in \mathcal{B}(\bar{\mathcal{X}};R,r)$. The construction of $\mathcal{P}r^i$ goes as follows. Let (L,k) be an object of $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)$, where $L \in \mathcal{X}$, $k \in \mathcal{B}(\bar{\mathcal{X}};R,r)$. We define $\mathcal{P}r^i((L,k)) = L$.

Next, we define the first-order part $\mathcal{P}r^i_1$ of $\mathcal{P}r^i$ on morphisms. This can be done by means of *Floer continuation maps*. Specifically, we need to define a filtration-preserving chain map

$$\mathcal{P}r^i_1 : \mathrm{CF}((L',k'),(L'',k'');\mathcal{D}_{(L',k'),(L'',k'')}) \to \mathrm{CF}(L',L'';\mathcal{D}_{L',L'',i}) \tag{3.24}$$

for every $L',L'' \in \mathcal{X}$, $k',k'' \in \mathcal{B}(\bar{\mathcal{X}};R,r)$, where $\mathcal{D}_{(L',k'),(L'',k'')}$ is the Floer data for the pair $((L',k'),(L'',k''))$ and $\mathcal{D}_{L',L'',i}$ is the one used in $\mathcal{F}uk(\mathcal{X};i)$. Assume first that $\bar{L}' \neq \bar{L}''$ (which means that $\bar{L}' \pitchfork \bar{L}''$). By construction, the Hamiltonian terms in both Floer data $\mathcal{D}_{(L',k'),(L'',k'')}$ and $\mathcal{D}_{L',L'',i}$ are identically 0 (so the two Floer data may differ only in their almost complex structures). Denote by $J(\mathcal{P}^{\mathrm{tot}};(L',k',L'',k''))$ the almost complex structure of $\mathcal{D}_{(L',k'),(L'',k'')}$ and by $J(i;L',L'')$ the one corresponding to $\mathcal{D}_{L',L'',i}$. Fix a generic homotopy J^{cont}_s, $s \in [0,1]$, between these two almost complex structures (of the type admissible in $\mathcal{B}(\bar{\mathcal{X}};R,r)$). We assume that J^{cont}_s coincides with $J(\mathcal{P}^{\mathrm{tot}};(L',k',L'',k''))$ near $s=0$ and with $J(i;L',L'')$ near $s=1$. Extend this homotopy to $s \in \mathbb{R}$ by keeping it constant with respect to the s-parameter outside $[0,1]$. Recall that, by construction, both Floer data $\mathcal{D}_{(L',k'),(L'',k'')}$ and $\mathcal{D}_{L',L'',i}$ have 0 Hamiltonian terms, hence $(0,J^{\mathrm{cont}}_s)$ defines a homotopy between the latter two Floer data. Standard Floer theory associates to this homotopy a filtration-preserving quasi-isomorphism, as claimed in (3.24), called *the Floer continuation map*. That $\mathcal{P}r^i_1$ preserves filtrations follows from standard arguments in Floer theory, using the assumption that the Hamiltonian terms in the preceding homotopy of Floer data are 0 for all times s.

For further use, we will add one more restriction on the definition of the continuation maps $\mathcal{P}r_1^i$. In case the two objects (L', k') and (L'', k'') correspond to the same perturbation data i, i.e., $k' = k'' = i$, we will take the homotopy J_s^{cont} to be constant. As a result, the continuation map

$$\mathcal{P}r_1^i : \mathrm{CF}((L', i), (L'', i); \mathcal{D}_{(L',i),(L'',i)}) \to \mathrm{CF}(L', L''; \mathcal{D}_{L',L'',i})$$

for such pairs will be the identity.

We now briefly address the case when $\overline{L}' = \overline{L}''$. Recall that in this case the Floer data on each of $\mathcal{D}_{(L',k'),(L'',k'')}$, $\mathcal{D}_{L',L'',i}$ consists of a Morse function and a Riemannian metric on $\overline{L}'$. In our model the corresponding CF's are just the Morse complexes on $\overline{L}'$ associated to these data. The map $\mathcal{P}r_1^i$ is now defined by means of standard Morse homology theory – it is just the continuation map between the two Morse complexes. The fact that $\mathcal{P}r_1^i$ is filtration preserving is automatic since, by definition, both Morse complexes $\mathrm{CF}((L', k'), (L'', k''); \mathcal{D}_{(L',k'),(L'',k'')})$ and $\mathrm{CF}(L', L''; \mathcal{D}_{L',L'',i})$ are concentrated at the same filtration level (which is a constant that depends on the difference between the primitives of the Liouville forms on L' and L'').

Similarly to the case $\overline{L}' \neq \overline{L}''$, here too we can arrange $\mathcal{P}r_1^i$ to be the identity map whenever $k' = k'' = i$. This can be done by taking the homotopy between the two Morse data to be constant, and the resulting Morse continuation map will then be the identity.

Next we define the second-order part of $\mathcal{P}r^i$. This will be a map

$$\mathcal{P}r_2^i : \mathrm{CF}((L_0, k_0), (L_1, k_1); \mathcal{D}_{\mathrm{tot}}) \otimes \mathrm{CF}((L_1, k_1), (L_2, k_2); \mathcal{D}_{\mathrm{tot}}) \to \mathrm{CF}(L_0, L_2; \mathcal{D}_i) \tag{3.25}$$

of cohomological degree -1. Here we have written $\mathcal{D}_{\mathrm{tot}}$ and $\mathcal{D}_i$ for the Floer data (for the corresponding pairs of Lagrangians) in the categories $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$ and $\mathcal{F}uk(\mathcal{X}; i)$, respectively.

Assume for simplicity that $\overline{L}_0$, $\overline{L}_1$, $\overline{L}_2$ are all distinct. To define the map (3.25) we will first need to introduce some new spaces of Floer-type polygons. Recall the homotopy J_s^{cont} from the definition of $\mathcal{P}r_1^i$ above. Below we will need a more precise notation, and from now on we will denote it by $J_{s,t}^{\mathrm{cont}}((L', k'), (L'', k''), (L', L'', i))$, where $s \in \mathbb{R}$ and $t \in [0, 1]$. (Recall that all our almost complex structures are possibly time-dependent, and we denote here by t the time parameter.)

Denote by S_3 the 3-punctured disk. Recall that S_3 has two "entry" strip-like ends $St_{0,1}^-$, $St_{1,2}^-$ and one "exit" strip-like end $St_{0,2}^+$. We order them in the clockwise direction, $St_{0,1}^-$, $St_{1,2}^-$, $St_{0,2}^+$, according to the punctures they correspond to. We denote by (s, t) the coordinates on each of these strip-like ends. Thus $(s, t) \in (-\infty, 0] \times [0, 1]$ for $St_{0,1}^-$, $St_{1,2}^-$, and $(s, t) \in [0, \infty) \times [0, 1]$ for $St_{0,2}^+$. We also fix a smooth positive decreasing function $A : (0, \delta) \to \mathbb{R}$ in a small neighborhood of 0 with $A(\tau) \to \infty$ as $\tau \to 0^+$.

Consider now a 1-parametric family $\mathcal{P}^\tau(\mathcal{P}\kappa) = (K^\tau(\mathcal{P}\kappa), J^\tau(\mathcal{P}\kappa))$, $\tau \in (0,1)$, of perturbation data on the 3-punctured disk S_3. We will write

$$\mathcal{P}^\tau_z(\mathcal{P}\kappa) = (K^\tau_z(\mathcal{P}\kappa), J^\tau_z(\mathcal{P}\kappa))$$

for the value of the perturbation data at the point $z \in S_3$. We will require the family $\mathcal{P}^\tau(\mathcal{P}\kappa)$, $\tau \in (0,1)$, to satisfy the following conditions:

(1) When $0 < \tau < \delta$, for every point in $St^-_{0,1}$ with coordinates $(s,t) \in (-\infty, 0] \times [0,1]$ we have

$$J^\tau_{s,t}(\mathcal{P}\kappa) = J^{\mathrm{cont}}_{s+A(\tau)+1,t}((L_0,k_0),(L_1,k_1),(L_0,L_1,i)).$$

In other words, when $s \in [-A(\tau)-1, -A(\tau)]$, $J^\tau_{s,-}(\mathcal{P}\kappa)$ coincides with the continuation homotopy J^{cont} after a suitable shift in the s-parameter. Note that $J^{\mathrm{cont}}_{s+A(\tau)+1,t}$ has been defined for all $s \in \mathbb{R}$.

(2) When $0 < \tau < \delta$, for every point in $St^-_{1,2}$ with coordinates $(s,t) \in (-\infty, 0] \times [0,1]$ we have

$$J^\tau_{s,t}(\mathcal{P}\kappa) = J^{\mathrm{cont}}_{s+A(\tau)+1,t}((L_1,k_1),(L_2,k_2),(L_1,L_2,i)).$$

(3) When $0 < \tau < \delta$, we require the perturbation data $\mathcal{P}^\tau(\mathcal{P}\kappa)$ to coincide with the one assigned by the perturbation data i to the triple (L_0,L_1,L_2) along the complement $S_3 \setminus (St^-_{0,1} \cup St^-_{1,2})$. Note that this requirement is compatible with the previous two conditions.

(4) When $1-\delta < \tau < 1$, for every point in $St^+_{0,2}$ with coordinates $(s,t) \in [0,\infty) \times [0,1]$ we have

$$J^\tau_{s,t}(\mathcal{P}\kappa) = J^{\mathrm{cont}}_{s-A(1-\tau),t}((L_0,k_0),(L_2,k_2),(L_0,L_2,i)).$$

(5) When $1-\delta < \tau < 1$, we require the perturbation data $\mathcal{P}^\tau(\mathcal{P}\kappa)$ to coincide with the one assigned by $\mathcal{P}^{\mathrm{tot}}$ to the triple $((L_0,k_0),(L_1,k_1),(L_2,k_2))$ along $S_3 \setminus St^+_{0,2}$. Again, this requirement is compatible with the previous one.

(6) For every $\tau \in [\delta, 1-\delta]$, the perturbation data $\mathcal{P}^\tau(\mathcal{P}\kappa)$ coincides with the Floer datum of the pair $((L_0,k_0),(L_1,k_1))$ (as assigned by $\mathcal{P}^{\mathrm{tot}}$) along $St^-_{0,1}$ outside some compact subset. We require the analogous condition to hold also with respect to $((L_1,k_1),(L_2,k_2))$ along $St^-_{1,2}$.

(7) For every $\tau \in [\delta, 1-\delta]$, the perturbation data $\mathcal{P}^\tau(\mathcal{P}\kappa)$ coincides with the Floer datum of the pair (L_0,L_2), as assigned by the perturbation data i, along $St^+_{0,2}$ outside some compact subset.

(8) For all $\tau \in (0,1)$, the almost complex structures $J^\tau(\mathcal{P}\kappa)$ and perturbation forms $K^\tau(\mathcal{P}\kappa)$ from $\mathcal{P}^\tau(\mathcal{P}\kappa)$ are all of the types and sizes admissible in the construction of $\mathcal{P}^{\mathrm{tot}}$. In particular, they should satisfy the inequality (3.23)

(for $d+1=3$) and the two conditions on K and J that are listed before that inequality on page 142.

(9) In case the three pairs (L_0,k_0), (L_1,k_1), (L_2,k_2) all correspond to the perturbation data i, i.e., $k_0=k_1=k_2$, we take the family $\{\mathcal{P}^\tau(\mathcal{P}r)\}$ to be constant with respect to τ, and moreover to coincide with the perturbation data i. Note that this is compatible with the rest of the conditions above, since earlier we required each of the homotopies J_s^{cont} that appear in points (1), (2), and (4) above to be constant (with respect to s).

Fix three intersection points $x\in\bar{L}_0\cap\bar{L}_1$, $y\in\bar{L}_1\cap\bar{L}_2$, $w\in\bar{L}_0\cap\bar{L}_2$. Denote by $\mathcal{M}^{\mathcal{P}r_2}(x,y,w;\{\mathcal{P}^\tau(\mathcal{P}r)\})$ the space of all pairs (η,u) where $\eta\in(0,1)$ and $u:S_3\to X$ solves the generalized Floer equation with respect to the perturbation data $\mathcal{P}^\eta(\mathcal{P}r)$, with Lagrangian boundary conditions prescribed by $\bar{L}_0,\bar{L}_1,\bar{L}_2$ and with asymptotics at the ends being x, y, and w.

By choosing the family $\{\mathcal{P}^\tau(\mathcal{P}r)\}$ to be generic, $\mathcal{M}^{\mathcal{P}r_2}(x,y,w;\{\mathcal{P}^\tau(\mathcal{P}r)\})$ is a smooth manifold of dimension $d(x,y,w)=|w|'-|x|'-|y|'+1$, where $|\cdot|'$ denotes cohomological degree. Moreover, standard arguments show that if $d(x,y,z)=0$ then $\mathcal{M}^{\mathcal{P}r_2}(x,y,z;\{\mathcal{P}^\tau(\mathcal{P}r)\})$ is compact, hence consists of finitely many points. We then define

$$\mathcal{P}r_2^i(x,y):=\sum_z \#\mathcal{M}^{\mathcal{P}r_2}(x,y,z;\{\mathcal{P}^\tau(\mathcal{P}r)\})z,$$

where z runs over all points in $\bar{L}_0\cap\bar{L}_2$ such that $|z|'=|x|'+|y|'-1$ (hence $d(x,y,z)=0$). As before, $\#\mathcal{M}^{\mathcal{P}r_2}(-)$ is the count (with values in $\mathbb{Z}_2$, or under additional assumptions in $\mathbf{k}$) of elements in the above space.

We claim that $\mathcal{P}r_1^i$ and $\mathcal{P}r_2^i$ form together an A_2-functor. This amounts to showing that $\mathcal{P}r_2^i$ satisfies the identity

$$\begin{aligned}&\mathcal{P}r_1^i(\mu_2^{\mathrm{tot}}(x,y))-\mu_2^{(i)}(\mathcal{P}r_1^i(x),\mathcal{P}r_1^i(y))\\&\quad=\epsilon_1\mu_1^{(i)}\mathcal{P}r_2^i(x,y)+\epsilon_2\mathcal{P}r_2^i(\mu_1^{\mathrm{tot}}(x),y)+\epsilon_3\mathcal{P}r_2^i(x,\mu_1^{\mathrm{tot}}(y)),\end{aligned}\tag{3.26}$$

where $\epsilon_1,\epsilon_2,\epsilon_3=\pm1$ are signs that depend on the degrees of x, y, and z (but as we work with $\mathbf{k}=\mathbb{Z}_2$, the precise value of these signs is irrelevant). Here we have denoted by μ^{tot} the A_∞-operations in the category $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)$ and by $\mu^{(i)}$ those from the category $\mathcal{F}uk(\mathcal{X};i)$.

The proof of (3.26) is based on standard arguments in Floer theory. Fix w such that $d(x,y,w)=1$. Then $\mathcal{M}^{\mathcal{P}r_2}(x,y,w;\{\mathcal{P}^\tau(\mathcal{P}r)\})$ is a 1-dimensional smooth manifold. Its compactification $\bar{\mathcal{M}}^{\mathcal{P}r_2}(x,y,w)$ is a compact 1-dimensional smooth manifold with boundary. The boundary points of $\partial\bar{\mathcal{M}}^{\mathcal{P}r_2}(x,y,w)$ consist of five types of broken trajectories, as depicted in Figure 3.9. The first two types correspond to $\tau\to1^-$ and $\tau\to0^+$, respectively, and the number of occurrences of each of them equals the coefficient of w in the first and second terms on the left-hand side of (3.26),

respectively. The other three types of broken trajectories, occurring at instances of time $0 < \tau_0 < 1$, correspond to standard breaking along strip-like ends. The coefficient of w in each of the terms on the right-hand side of (3.26) equals the number of occurrences of broken trajectories of the corresponding type. The identity (3.26) now follows (with appropriate sings ϵ_i) since the signed number of boundary points in $\bar{\mathcal{M}}^{\mathcal{P}r_2}(x, y, w)$ must be 0. This concludes the construction of the A_2-functor $\mathcal{P}r^i$ for every $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$.

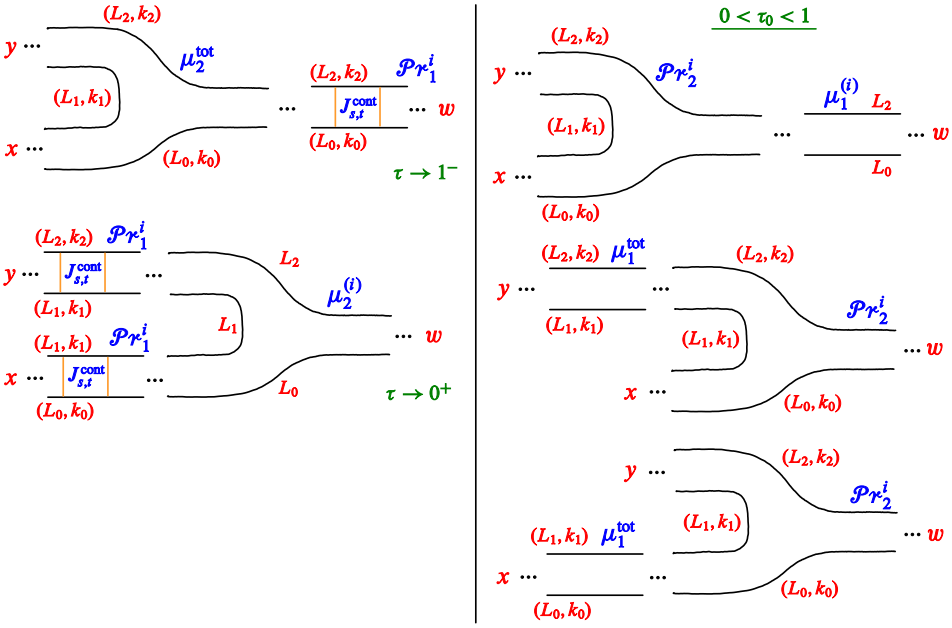


Figure 3.9. The five possible types of boundary points of $\partial\bar{\mathcal{M}}^{\mathcal{P}r_2}(x, y, w)$.

We now claim that $\mathcal{P}r^i$ preserves filtrations. We have already proved earlier that $\mathcal{P}r_1^i$ is filtration preserving, so it remains to deal with the second component $\mathcal{P}r_2^i$. The fact that $\mathcal{P}r_2^i$ is filtration preserving follows from condition (8) on page 145. Indeed, due to inequality (3.13), the proof from pages 127–131 extends with minor modifications to show that, for every $(\eta, u) \in \mathcal{M}^{\mathcal{P}r_2}(x, y, z; \{\mathcal{P}^\tau(\mathcal{P}r)\})$, the symplectic area of u satisfies $A(u) > 0$.

We will soon apply Proposition 3.20 with $\mathcal{A}^i = \mathcal{F}uk(\mathcal{X}; i)$, $i \in \mathcal{I} := \mathcal{B}(\bar{\mathcal{X}}; R, r)$, $\mathcal{A}^{\mathrm{tot}} = \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$. In order to do so, we still need to show several other properties of the A_2-functors $\mathcal{P}r^i$, as required by Proposition 3.20.

The first one is that $\mathcal{P}r^i$ is strictly A_1-unital. This follows immediately from Morse theory, since the continuation maps in Morse theory send the unique local

maximum of the Morse function on (L,k) to the corresponding one on (L,i). Therefore $\mathcal{P}r^i_1$ sends strict units to strict units.

The next property is that $\mathcal{P}r^i \circ \mathcal{J}^i = \mathbb{1}_{\mathcal{F}uk(\mathcal{X};i)}$ as A_2-functors. This will follow from the following two statements:

(1) For all pairs of the type $((L',i),(L'',i))$ we have $\mathcal{P}r^i_1 = \mathbb{1}$.

(2) For any three objects of the type $(L_0,i),(L_1,i),(L_2,i)$ we have $\mathcal{P}r^i_2 = 0$.

The first statement has already been proved earlier. The second one follows from condition (9) in the definition of the family $\{\mathcal{P}^\tau(\mathcal{P}r)\}$, namely the requirement that this family is constant (in τ) and coincides with i for all τ. This implies that $\mathcal{M}^{\mathcal{P}r_2}(x,y,z)$ is the same space as the one defining the operation μ_2 (in the category $\mathcal{F}uk(\mathcal{X};i)$). A transversality/dimension argument now shows that whenever $|z|' = |x|' + |y|' - 1$ we have $\mathcal{M}^{\mathcal{P}r_2}(x,y,z) = \emptyset$. This proves the second statement and concludes the proof that $\mathcal{P}r^i \circ \mathcal{J}^i = \mathbb{1}_{\mathcal{F}uk(\mathcal{X};i)}$ as A_2-functors.

Finally, we claim that the A_2-functor $\mathcal{P}r^i$ has the property described just before the statement of Proposition 3.20. Namely, for every object (L,k) of $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)$ there exists a strict isomorphism $u^i_{(L,k)} \in \hom^{\le 0}_{\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)}((L,k),(L,i))$. Indeed, we can take $u^i_{(L,k)}$ to be the unique critical point of index $n = \dim_{\mathbb{C}} X$ for the Morse function in the Floer datum of the pair $((L,k),(L,i))$. Standard arguments in Morse theory then show that $u^i_{(L,k)}$ is a strict isomorphism.

We are now in a position to apply Proposition 3.20 with $\mathcal{A}^i = \mathcal{F}uk(\mathcal{X};i)$, $i \in \mathcal{I} := \mathcal{B}(\overline{\mathcal{X}};R,r)$, $\mathcal{A}^{\mathrm{tot}} = \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)$. By that proposition we obtain the structure of a coherent system of filtered A_∞-categories on the family $\{\mathcal{F}uk(\mathcal{X};i)\}_{i\in\mathcal{B}(\overline{\mathcal{X}};R,r)}$. Note that this holds for every $r \le R$, hence in particular also for $r = R$. Note also that $\mathcal{B}(\overline{\mathcal{X}};R,r) \subset \mathcal{B}(\overline{\mathcal{X}};R,R)$, and we can arrange our choices (of $\mathcal{P}^{\mathrm{tot}}$) such that the coherent system over the larger base $\mathcal{B}(\overline{\mathcal{X}};R,R)$ restricts to the one over the smaller base $\mathcal{B}(\overline{\mathcal{X}};R,r)$.

The above construction gives us many coherent systems. Namely, one coherent system over each base $\mathcal{B}(\overline{\mathcal{X}};R,R)$ for all $0 < R \le R_{\overline{\mathcal{X}}}$. We denote the comparison functors of the coherent system over $\mathcal{B}(\overline{\mathcal{X}};R,R)$ by $\mathcal{F}^{j,i}_R : \mathcal{F}uk(\mathcal{X};i) \to \mathcal{F}uk(\mathcal{X};j)$ for every $i,j \in \mathcal{B}(\overline{\mathcal{X}};R,R)$. Similarly, we denote by T^{i_2,i_1,i_0}_R the natural transformations of this system (relating $\mathcal{F}^{i_2,i_1}_R \circ \mathcal{F}^{i_1,i_0}_R$ to $\mathcal{F}^{i_2,i_0}_R$).

Remark 3.21. Ideally, we would have liked to construct one total category $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};\mathcal{B}(\overline{\mathcal{X}}))$ over the entire base $\mathcal{B}(\overline{\mathcal{X}})$. Unfortunately this is not so straightforward to achieve, at least not with our construction of the total categories. The difficulty has to do with establishing a set of perturbation data $\mathcal{P}^{\mathrm{tot}}$ (over the entire base $\mathcal{B}(\overline{\mathcal{X}})$, and in fact even over subspaces of it of the form $\mathcal{B}(\overline{\mathcal{X}};R_1,R_1) \cup \mathcal{B}(\overline{\mathcal{X}};R_2,R_2)$) that is both consistent (with respect to splitting/gluing) and at the same time yields filtration-preserving operations μ_d.

To understand better the difficulty, consider the case of the union $\mathcal{B}(\bar{\mathcal{X}}; R_1, R_1) \cup \mathcal{B}(\bar{\mathcal{X}}; R_2, R_2)$, where $R_2 < R_1 \leq R_{\bar{\mathcal{X}}}$. Let (i_0, i_1, i_2, i_3) be a tuple of perturbation data with

$$i_1, i_3 \in \mathcal{B}(\bar{\mathcal{X}}; R_1, R_1) \cap \mathcal{B}(\bar{\mathcal{X}}; R_2, R_2) = \mathcal{B}(\bar{\mathcal{X}}; R_1, R_2),$$
$$i_2 \in \mathcal{B}(\bar{\mathcal{X}}; R_1, R_1) \setminus \mathcal{B}(\bar{\mathcal{X}}; R_2, R_2), \quad i_0 \in \mathcal{B}(\bar{\mathcal{X}}; R_2, R_2) \setminus \mathcal{B}(\bar{\mathcal{X}}; R_1, R_1).$$

For some choices of i_0, i_1, i_2, i_3, we might have

$$R^{(i_0,i_1,i_2,i_3)} = R_2, \quad R^{(i_1,i_2,i_3)} = R_1, \quad R^{(i_0,i_1,i_3)} = R_2,$$
$$r^{(i_0,i_1,i_2,i_3)} = R_1, \quad r^{(i_1,i_2,i_3)} = R_1, \quad r^{(i_0,i_1,i_3)} = R_2,$$

where $R^{(\cdots)}$ and $r^{(\cdots)}$ are defined in (3.21). Now recall that in order to obtain a choice of perturbation data that preserves filtrations we used inequalities (3.22) and (3.23). Since $R_2 < R_1$, in our case we have $R^{(i_0,i_1,i_2,i_3)} < r^{(i_0,i_1,i_2,i_3)}$, so inequality (3.22) does not hold. There are also problems regarding the consistency of the perturbation data with respect to splitting and gluing. Indeed, a Floer polygon labeled by $((L_0, i_0), (L_1, i_1), (L_2, i_2), (L_3, i_3))$ can split into two polygons labeled by $((L_1, i_1), (L_2, i_2), (L_3, i_3))$ and by $((L_0, i_0), (L_1, i_1), (L_3, i_3))$. And vice versa, pairs of polygons of the latter type can be glued into polygons of the former type (assuming obvious matching assumptions). In our case we have $r^{(i_1,i_2,i_3)} > r^{(i_0,i_1,i_3)}$, so in order to achieve the consistency of the perturbation data while at the same time having filtration-preserving μ_d's, we would need to decrease the size of the perturbation forms of $\mathcal{P}^{\mathrm{tot}}((L_1, i_1), (L_2, i_2), (L_3, i_3))$ to at most $r^{(i_0,i_1,i_3)}$. In turn, this might not be compatible with the size of the perturbation forms of the data i_2 (which may be of size R_1). Similar problems arise also with the behavior of the measurement $R^{(i_0,\ldots,i_d)}$ with respect to splitting/gluing when the indices i_k are spread in spaces $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ with different values of R.

Summing up, while our construction for total categories works over bases like $\mathcal{B}(\bar{\mathcal{X}}; R, R)$, it is not clear whether the construction can be extended over a base of the type $\mathcal{B}(\bar{\mathcal{X}}; R_1, R_1) \cup \mathcal{B}(\bar{\mathcal{X}}; R_2, R_2)$.

In view of the above, instead of constructing a total category over $\mathcal{B}(\bar{\mathcal{X}})$ we will simply try to extend the various coherent systems (coming from the total categories) over each $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ to coherent systems over unions of such subspaces.

We proceed now with the extension of the system of comparison functors to spaces beyond the type $\mathcal{B}(\bar{\mathcal{X}}; R, R)$. Recall that if $R' \leq R \leq R_{\bar{\mathcal{X}}}$ then $\mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R') = \mathcal{B}(\bar{\mathcal{X}}; R, R')$. We claim that the construction of the total categories above and the coherent systems resulting from them can be carried out in such a way that the following holds for every $R' \leq R \leq R_{\bar{\mathcal{X}}}$: the two coherent systems, one over $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ and the other over $\mathcal{B}(\bar{\mathcal{X}}; R', R')$, coincide over the overlap $\mathcal{B}(\bar{\mathcal{X}}; R, R')$. In other words, the functors $\mathcal{F}_R^{j,i}$ and $\mathcal{F}_{R'}^{j,i}$ coincide for every $i, j \in \mathcal{B}(\bar{\mathcal{X}}; R, R')$, and similarly for the natural transformations $T_R^{i_2,i_1,i_0}$, $T_{R'}^{i_2,i_1,i_0}$.

The proof is based on three steps. The first one is that we can choose the perturbation data for the total categories in such a way that the categories $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R)$ and $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R', R')$ coincide along the overlaps $\mathcal{B}(\overline{\mathcal{X}}; R, R')$ (and are equal along that overlap to $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$). The way we chose the perturbation data $\mathcal{P}^{\mathrm{tot}}$ makes this possible, since in case $i_0, \dots, i_d \in \mathcal{B}(\overline{\mathcal{X}}; R, R) \cap \mathcal{B}(\overline{\mathcal{X}}; R', R')$ the definitions of the parameters $R^{(i_0,\dots,i_d)}$ and $r^{(i_0,\dots,i_d)}$ are independent of whether we view the indices i_k as elements of $\mathcal{B}(\overline{\mathcal{X}}; R, R)$ or of $\mathcal{B}(\overline{\mathcal{X}}; R', R')$.

The second step has to do with the A_2-functors $\mathcal{P}r^i$. To keep track of the domain of these functors, we will temporarily add a subscript R to their notation:

$$\mathcal{P}r^i_R : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R) \to \mathcal{F}uk(\mathcal{X}; i), \quad i \in \mathcal{B}(\overline{\mathcal{X}}; R, R);$$

we also have

$$\mathcal{P}r^i_{R'} : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R) \to \mathcal{F}uk(\mathcal{X}; i), \quad i \in \mathcal{B}(\overline{\mathcal{X}}; R', R').$$

Recall that $\mathcal{B}(\overline{\mathcal{X}}; R, R) \cap \mathcal{B}(\overline{\mathcal{X}}; R', R') = \mathcal{B}(\overline{\mathcal{X}}; R, R')$ and that $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$ is a subcategory of both $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R)$ and $\mathcal{B}(\overline{\mathcal{X}}; R', R')$, which are the domains of $\mathcal{P}r^i_R$ and $\mathcal{P}r^i_{R'}$, respectively. We need to show that for every $i \in \mathcal{B}(\overline{\mathcal{X}}; R, R')$ the two A_2-functors $\mathcal{P}r^i_R$ and $\mathcal{P}r^i_{R'}$ coincide when restricted to the subcategory $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$ of their respective domains. This can be achieved by arranging the 1-parametric families of the perturbation data $\{\mathcal{P}^\tau(\mathcal{P}r)\}_{\tau\in(0,1)}$ for $\mathcal{B}(\overline{\mathcal{X}}; R, R)$ and $\mathcal{B}(\overline{\mathcal{X}}; R', R')$ to coincide over their intersection.

The third and last step concerns the extensions of the preceding A_2-functors to A_∞-functors, as stated in Proposition 3.20. Here we need to show that the A_∞-extension $\mathcal{Q}^i_R$ of $\mathcal{P}r^i_R$ can be arranged to coincide with the A_∞-extension $\mathcal{Q}^i_{R'}$ of $\mathcal{P}r^i_{R'}$ over the common subcategory of their domains $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$ whenever i is in $\mathcal{B}(\overline{\mathcal{X}}; R, R')$.

This can be proved by means of the algebraic Lemmas 3.18 and 3.19 and Proposition 3.20. Note that for $i \in \mathcal{B}(\overline{\mathcal{X}}; R, R')$, the inclusion of $\mathcal{F}uk(\mathcal{X}; i)$ into any of the categories $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$, $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R)$, and $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R', R')$ is a filtered quasi-equivalence.

We now claim that the four maps

$$\begin{aligned}
&\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R), \mathcal{F}uk(i); \mathcal{P}r^i_R, \mathcal{P}r^i_R\big) \to \mathrm{PHH}(\mathcal{F}uk(i), \mathcal{F}uk(i); \mathbb{1}, \mathbb{1}),\\
&\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R', R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R'}, \mathcal{P}r^i_{R'}\big) \to \mathrm{PHH}(\mathcal{F}uk(i), \mathcal{F}uk(i); \mathbb{1}, \mathbb{1}),\\
&\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R), \mathcal{F}uk(i); \mathcal{P}r^i_R, \mathcal{P}r^i_R\big)\\
&\qquad\qquad \to \mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R,R'}, \mathcal{P}r^i_{R,R'}\big),\\
&\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R', R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R'}, \mathcal{P}r^i_{R'}\big)\\
&\qquad\qquad \to \mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R,R'}, \mathcal{P}r^i_{R,R'}\big),
\end{aligned}$$

induced by the obvious restrictions, are all isomorphisms of persistence modules. Here we have omitted "$\mathcal{X}$" from all the $\mathcal{F}uk$-categories in an attempt to keep the formulas short. The A_2-functor $\mathcal{P}r^i_{R,R'} : \mathcal{F}uk^{\mathrm{tot}}(R, R') \to \mathcal{F}uk(i)$ appearing in the third and fourth maps above is just $\mathcal{P}r^i_{R,R'} := \mathcal{P}r^i_R|_{\mathcal{F}uk^{\mathrm{tot}}(R,R')} = \mathcal{P}r^i_{R'}|_{\mathcal{F}uk^{\mathrm{tot}}(R,R')}$.

Indeed, that the first two maps are isomorphisms follows from Lemma 3.18. For the third map, note that, by Lemma 3.18, for every $R' \leq R$ the map

$$\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R,R'}, \mathcal{P}r^i_{R,R'}\big) \to \mathrm{PHH}(\mathcal{F}uk(i), \mathcal{F}uk(i); \mathbb{1}, \mathbb{1}) \tag{3.27}$$

(also induced by restriction) is an isomorphism of persistence modules. Now, the first of the above four maps factors as the composition of the third map with the map in (3.27). Since the latter map and the first map are both isomorphisms, it follows that the same holds for the third map. The proof that the fourth map is an isomorphism is similar.

We proceed now by extending $\mathcal{P}r^i_R|_{\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,R')} = \mathcal{P}r^i_{R'}|_{\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,R')}$ to an A_∞-functor

$$\mathcal{Q}^i_{R,R'} : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R') \to \mathcal{F}uk(\mathcal{X}; i).$$

Next we extend $\mathcal{Q}^i_{R,R'}$ twice more: once to an A_∞-functor $\mathcal{Q}^i_R : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R) \to \mathcal{F}uk(\mathcal{X}; i)$ and another time to an A_∞-functor $\mathcal{Q}^i_{R'} : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R', R') \to \mathcal{F}uk(\mathcal{X}; i)$. This concludes the third step.

Finally, it is possible to show that the systems of natural transformations T^{i_2,i_1,i_0}_R and $T^{i_2,i_1,i_0}_{R'}$ can be chosen to agree over the intersection $\mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R')$. This can be done by arguments similar to those above (based on the construction of the natural transformations from the proof of Proposition 3.20).

Recall that the comparison functors $\mathcal{F}^{j,i}_R$ in the coherent system over $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ are given by $\mathcal{F}^{j,i}_R = \mathcal{Q}^j_R \circ \mathcal{J}^i_R$, where $\mathcal{J}^i_R : \mathcal{F}uk(\mathcal{X}; i) \to \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R)$ is the inclusion functor. The comparison functors over $\mathcal{B}(\bar{\mathcal{X}}; R', R')$ and $\mathcal{B}(\bar{\mathcal{X}}; R, R')$ have essentially the same expressions.

Before we proceed, let us summarize what we have proven so far.

Proposition 3.22. *The space of perturbation data $\mathcal{B}(\bar{\mathcal{X}})$ can be written as a union*

$$\mathcal{B}(\bar{\mathcal{X}}) = \bigcup_{R \leq R_{\bar{\mathcal{X}}}} \mathcal{B}(\bar{\mathcal{X}}; R, R)$$

of mutually overlapping subspaces $\mathcal{B}(\bar{\mathcal{X}}; R, R)$. The overlap between any two of these subspaces is given by $\mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R') = \mathcal{B}(\bar{\mathcal{X}}; R, R')$ for every $R' \leq R$.

The family of Fukaya categories $\{\mathcal{F}uk(\mathcal{X}; i)\}_{i \in \mathcal{B}(\bar{\mathcal{X}};R,R)}$ over each subspace $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ can be endowed with a coherent system of comparison A_∞-functors that are filtered quasi-equivalences and are strictly A_1-unital. Moreover, for every

R, R', there exist coherent systems over $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ and over $\mathcal{B}(\bar{\mathcal{X}}; R', R')$ that agree along the overlap $\mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R')$.

Finally, all the comparison functors in the various coherent systems above act as the identity maps on the objects of the respective categories. The first-order terms of the comparison functors induce the canonical continuation maps in Floer homology.

The structure of the subspaces $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ and their overlaps is schematically depicted in Figure 3.10.

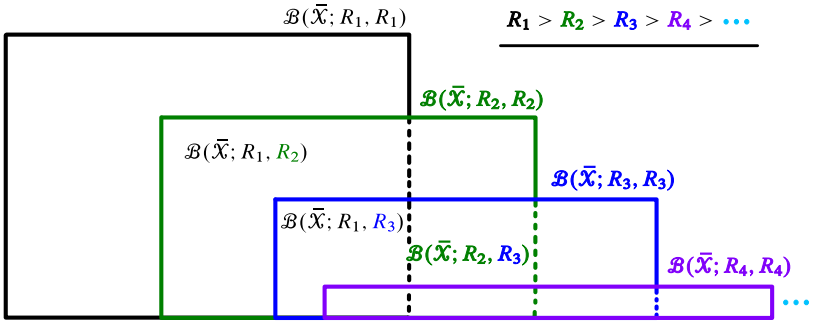


Figure 3.10. The subspaces $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ and their overlaps.

Remark 3.23. The type of structure described in Proposition 3.22 can be defined abstractly for any collection of A_∞-categories (filtered or not) $\{\mathcal{A}^i\}_{i\in\mathcal{B}}$ over a base $\mathcal{B}$. An appropriate name for such a structure could be *a collection of coherent systems with overlaps*. We will not pursue this direction further in this memoir and will now proceed with our Fukaya categories.

We are ready now to describe the functors $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ claimed in Theorem 3.12. To simplify the notation, we continue to denote perturbation data by indices like i instead of $\mathcal{P}$, and the comparison functors will be denoted by $\mathcal{F}^{j,i}$ instead of $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$. Note that the construction below is purely formal and can be applied to any family of A_∞-categories endowed with a collection of coherent systems with overlaps (as in Remark 3.23).

For every $R', R'' \leq R_{\bar{\mathcal{X}}}$ we choose a base point

$$l(R', R'') \in \mathcal{B}(\bar{\mathcal{X}}; \max\{R', R''\}, \min\{R', R''\}).$$

Given $i \in \mathcal{B}(\bar{\mathcal{X}})$, define $R^{(i)} := \sup\{\widetilde{R} \mid \widetilde{R} \leq R_{\bar{\mathcal{X}}} \text{ and } i \in \mathcal{B}(\bar{\mathcal{X}}; \widetilde{R}, \widetilde{R})\}$, which is a special case of the first parameter in (3.21).

Let $i, j \in \mathcal{B}(\bar{\mathcal{X}})$. If $i = j$, define $\mathcal{F}^{i,j} := \mathbb{1}_{\mathcal{F}uk(\mathcal{X};i)}$. If $i \neq j$, set $l_{i,j} := l(R^{(i)}, R^{(j)})$ and define

$$\mathcal{F}^{j,i} := \mathcal{F}^{j,l_{i,j}}_{R^{(j)}} \circ \mathcal{F}^{l_{i,j},i}_{R^{(i)}} : \mathcal{F}uk(\mathcal{X}; i) \to \mathcal{F}uk(\mathcal{X}; j). \tag{3.28}$$

It is easy to see that the functors $\mathcal{F}^{j,i}$ satisfy all the properties claimed in Theorem 3.12.

Remark 3.24. The system of functors $\mathcal{F}^{j,i}$ defined by (3.28) is not canonical. The construction uses many different choices at different stages. However, one can show that the dependence on the choices of the base points $l_{R',R''}$ is somewhat controlled. If one replaces the base points $l_{R',R''}$ by a different set of choices $l'_{R',R''}$, then the resulting system of comparison functors $\mathcal{F}'^{j,i}$ will be naturally quasi-isomorphic to the system $\mathcal{F}^{j,i}$ by natural quasi-isomorphisms that preserve filtrations.

The applications in this memoir do not require the comparison functors or the natural transformations between them to be canonical. In fact, in what follows we just need to know that the different Fukaya categories $\mathcal{F}uk(\mathcal{X};i)$ are filtered quasi-equivalent.

It remains to address the last statement of Theorem 3.12, concerning the collection of Lagrangians $\bar{\mathcal{X}}'$. This follows immediately from our construction. Indeed, let $\bar{\mathcal{X}}'$ be a collection of Lagrangians with $\bar{\mathcal{X}}' \supset \bar{\mathcal{X}}$, and assume that $\bar{\mathcal{X}}'$ satisfies all the conditions from the beginning of Section 3.2.1. We can use the same symplectic embedding ϕ from (3.10) but with a smaller radius $R_{\bar{\mathcal{X}}'} \leq R_{\bar{\mathcal{X}}}$. Clearly $\mathcal{B}(\bar{\mathcal{X}}')|_{\bar{\mathcal{X}}} \subset \mathcal{B}(\bar{\mathcal{X}})$.

This completes the proof of Theorem 3.12. ∎

3.4 Proofs of the main symplectic applications

The first subsection is dedicated to the proofs of Theorems 3.1 and 3.4 and the second to the proof of Corollary 3.7.

3.4.1 Proofs of Theorems 3.1 and 3.4

We begin with Theorem 3.1, and then proceed to the proof of Theorem 3.4 in Section 3.4.1.2 on page 156.

3.4.1.1 Proof of Theorem 3.1. Recall from Theorem 3.12 that there are choices of Floer data and perturbation data $\mathcal{P}$ such that the resulting Fukaya category $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ is filtered, strictly unital, and, without filtrations, quasi-equivalent to the subcategory $\mathcal{F}uk(\mathcal{X})$ of $\mathcal{F}uk(X)$ whose collection of objects is $\mathcal{X}$ (see Section 3.1 for the notation).

We will now apply the discussion in Section 3.2.1 to the filtered A_∞-category $\mathcal{F}uk(\mathcal{X};\mathcal{P})$. The first step is to discuss the shift functor. There is an obvious shift functor on the category $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ which acts on objects by $\Sigma^r L = (\bar{L}, h_L + r, \theta_L)$ (see Section 3.1). This action on objects induces an A_∞-shift functor on $\mathcal{F}uk(\mathcal{X};\mathcal{P})$,

because the perturbation data $\mathcal{P}$ depends only on the geometric part of the marked Lagrangians. It is easy to see that this shift functor on $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ is compatible with the shift functor defined on $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$ as in Section 3.2.1: for a module $\mathcal{M} \in \mathrm{Fmod}(\mathcal{A})$, the filtered module $\Sigma^r\mathcal{M}$ is defined by $(\Sigma^r\mathcal{M})^{\leq\alpha}(N) = \mathcal{M}^{\alpha-r}(N)$, with the same μ_d-operations as $\mathcal{M}$. Indeed,

$$\Sigma^r(\mathcal{Y}(L)) = \mathcal{Y}(\Sigma^r L).$$

We now define

$$\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P}) = H^0(\mathcal{F}uk(\mathcal{X};\mathcal{P})^\nabla).$$

Recall that $\mathcal{F}uk(\mathcal{X};\mathcal{P})^\nabla$ is constructed by first considering the category $\mathcal{F}uk(\mathcal{X};\mathcal{P})^\#$ consisting of the triangulated completion of the Yoneda modules of the elements in $\mathcal{X}$. Triangles are understood here to be of the form $\mathcal{M} \xrightarrow{f} \mathcal{N} \to \mathrm{Cone}(f)$, where the cone construction is in the sense of filtered A_∞-modules and the morphism f preserves filtration.

The category $\mathcal{F}uk(\mathcal{X};\mathcal{P})^\nabla$ is the full subcategory of $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$ that contains $\mathcal{F}uk(\mathcal{X};\mathcal{P})^\#$, as well as all the modules in $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$, together with all their shifts and translates that are r-quasi-isomorphic to the objects in $\mathcal{F}uk(\mathcal{X};\mathcal{P})^\#$ for some $r \geq 0$. From the discussion above concerning the shift functor, it follows immediately that $\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P})$ is indeed a TPC.

Remark 3.25. More explicitly, to obtain $\mathcal{F}uk(\mathcal{X};\mathcal{P})^\nabla$ from $\mathcal{F}uk(\mathcal{X};\mathcal{P})^\#$, we first add to the Yoneda modules, together with all possible shifts and translates, the r-acyclic modules for all $r \geq 0$; then we take the triangulated completion of all of them.

We now start the proof by discussing the independence of $\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P})$ of the perturbation data $\mathcal{P}$, up to TPC equivalence. The argument is a direct consequence of the system of comparison functors for the categories $\mathcal{F}uk(\mathcal{X};\mathcal{P})$, as given in Theorem 3.12. Indeed, for two choices of admissible perturbation data $\mathcal{P}_1$ and $\mathcal{P}_2$, we have a filtered functor

$$\mathcal{F}^{1,2} : \mathcal{F}uk(\mathcal{X},\mathcal{P}_1) \to \mathcal{F}uk(\mathcal{X},\mathcal{P}_2).$$

The existence of the natural transformations that compare the compositions $\mathcal{F}^{1,2} \circ \mathcal{F}^{2,1}$, $\mathcal{F}^{2,1} \circ \mathcal{F}^{1,2}$ with the respective identities implies that the associated homological functor is full and faithful, in the sense that it induces an isomorphism of persistence modules $H(\mathcal{F}^{1,2}) : H^0(\hom^{\leq r}(X,Y)) \to H^0(\hom^{\leq r}(X,Y))$ for every $r \in \mathbb{R}$ and every two objects X, Y of $\mathcal{F}uk(\mathcal{X},\mathcal{P}_1)$. As in the unfiltered case, a consequence of the existence of these functors and the natural transformation relating them is that the pullback of filtered modules $[\mathcal{F}^{1,2}]^* : \mathrm{Fmod}(\mathcal{F}uk(\mathcal{X},\mathcal{P}_2)) \to \mathrm{Fmod}(\mathcal{F}uk(\mathcal{X},\mathcal{P}_1))$ sends each Yoneda module $\mathcal{Y}_{\mathcal{P}_2}(L)$ to a module 0-quasi-isomorphic to $\mathcal{Y}_{\mathcal{P}_1}(L)$ (in the sense that the two modules are related by a morphism that induces a 0-isomorphism in the homological category). The standard properties of the pullback of A_∞-modules

imply that $[\mathcal{F}^{1,2}]^*$ respects the triangulated structure with respect to 0-weight triangles as well as shift functors. We deduce that this pullback also sends r-isomorphisms to r-isomorphisms, and thus sends $\mathcal{F}uk(\mathcal{X},\mathcal{P}_2)^\nabla$ to $\mathcal{F}uk(\mathcal{X},\mathcal{P}_1)^\nabla$. Finally, using the relevant natural transformations, we deduce that $H[\mathcal{F}^{1,2}]^*$ is an equivalence of TPCs.

We will denote the resulting TPC by $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ in place of $\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P})$, except where this may lead to confusion.

We now continue with points (i), (ii), and (iii) of Theorem 3.1. The first point is immediate because $\hom_{\mathcal{C}\mathcal{F}uk(\mathcal{X})}(L,L') \cong \mathrm{HF}(L,L')$ as persistence modules, with the persistence structure on $\mathrm{HF}(L,L')$ as described in Section 3.2.2.

The second point of the theorem claims that the ∞-level of $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ is equivalent to $D\mathcal{F}uk(\mathcal{X})$. The argument is the following. First, it is well known [53] that an alternative, equivalent model for $D\mathcal{F}uk(\mathcal{X})$ is given by the homological category of twisted complexes, $H^0(\mathrm{Tw}(\mathcal{F}uk(\mathcal{X})))$. Here $\mathcal{F}uk(\mathcal{X})$ is any A_∞-category that represents the Fukaya category with objects the Lagrangians in $\mathcal{X}$. Thus we can take in its place $\mathcal{F}uk_{uf}(\mathcal{X};\mathcal{P})$, where the notation $(-)_{uf}$ means that we neglect the filtration. A variant of Lemma 2.96 applies also to our A_∞-setting, and it implies that each twisted complex in $\mathrm{Tw}(\mathcal{F}uk_{uf}(\mathcal{X};\mathcal{P}))$ can be viewed as a filtered twisted complex in $\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$ whose filtration is forgotten. Filtered twisted complexes are discussed in Section 3.2.1 (and in more detail in the dg case, in Section 2.5.1). By passing to homology, this means that $[H^0\,\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))]_\infty \cong H^0\,\mathrm{Tw}(\mathcal{F}uk_{uf}(\mathcal{X};\mathcal{P})) \cong D\mathcal{F}uk(\mathcal{X})$. In turn, $\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$ is easily seen to be equivalent to $\mathcal{F}uk(\mathcal{X};\mathcal{P})^\#$ as TPC categories. Thus,

$$[\mathcal{F}uk(\mathcal{X};\mathcal{P})^\nabla]_\infty \cong [\mathcal{F}uk(\mathcal{X};\mathcal{P})^\#]_\infty \cong [H^0\,\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))]_\infty \cong D\mathcal{F}uk(\mathcal{X}).$$

We now turn to the third point of Theorem 3.1. The argument here makes again essential use of the systems of comparison functors provided by Theorem 3.12, which allows us to extend perturbation data from one set $\mathcal{X}$ to a larger one, $\mathcal{X}'$.

Let N be a marked Lagrangian that is in general position with respect to the family $\overline{\mathcal{X}}$. We add N, as well as all its shifts and translates, to the family $\mathcal{X}$, thus obtaining $\mathcal{X}'$. By Theorem 3.12 and the invariance of $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ relative to perturbation data, we can now assume that the perturbation data $\mathcal{P}$ extends to a new one $\mathcal{P}'$ that defines the filtered A_∞-category $\mathcal{F}uk(\mathcal{X}',\mathcal{P}')$. Finally, we can extend $\mathcal{P}'$ to a perturbation data that defines an A_∞-category which is equivalent to $\mathcal{F}uk(X)$ if filtrations are forgotten. At this point, it is easier to pursue the argument using twisted complexes. As before, these will be of two types: filtered, those from $\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$, and unfiltered, those belonging to $\mathrm{Tw}(\mathcal{F}uk_{uf}(\mathcal{X}';\mathcal{P}'))$. The assumption that $\mathcal{X}$ generates $D\mathcal{F}uk(X)$ implies that there is an unfiltered quasi-isomorphism of twisted complexes $\phi : N \to C$, where C is a twisted module involving only elements of $\mathcal{X}$. Both N and C have structures of filtered twisted complexes, but the morphism ϕ, a priori, does not see the filtration.

However, using an argument similar to that in Lemma 2.96, we deduce that (after possibly shifting up N) we can view ϕ as a filtration-preserving morphism in $\hom_{\mathrm{Tw}(\mathcal{Fuk}(\mathcal{X}';\mathcal{P}'))}$. The mapping cone K of ϕ is acyclic as a twisted module (forgetting the filtration) because ϕ is a quasi-isomorphism. In other words, the identity 1_K of K is homologous to 0 in $\hom_{\mathrm{Tw}(\mathcal{Fuk}_{uf}(\mathcal{X}';\mathcal{P}'))}(K,K)$. Using again a reasoning similar to that in Lemma 2.96, we deduce that 1_K vanishes in

$$H^0\big(\hom^{\leq r}_{\mathrm{Tw}(\mathcal{Fuk}(\mathcal{X}';\mathcal{P}'))}(K,K)\big)$$

for some $r \geq 0$. This means that ϕ is an r-isomorphism. We can now reformulate the result in terms of modules, and we deduce that $j^*\mathcal{Y}(N)$ is r-isomorphic to a filtered module from $\mathcal{Fuk}(\mathcal{X},\mathcal{P})^{\#}$, where $j : \mathcal{Fuk}(\mathcal{X},\mathcal{P}) \to \mathcal{Fuk}(\mathcal{X}',\mathcal{P}')$ is the inclusion. This concludes the proof. ∎

3.4.1.2 Proof of Theorem 3.4. Theorem 3.12 implies that for every two admissible perturbation data $\mathcal{P}_1$ and $\mathcal{P}_2$ there exists a functor

$$[\mathcal{F}^{1,2}]^* : \mathcal{Fuk}(\mathcal{X},\mathcal{P}_2)^{\nabla} \to \mathcal{Fuk}(\mathcal{X},\mathcal{P}_1)^{\nabla}$$

that induces an equivalence of TPCs in homology and is the identity map on the objects from $\mathcal{X}$. Given that $\mathcal{C}\mathcal{Fuk}(\mathcal{X};\mathcal{P}_i) = H^0(\mathcal{Fuk}(\mathcal{X},\mathcal{P}_i)^{\nabla})$, this means that the pseudometric $D^{\mathcal{F}}$ defined on $\mathcal{X}$ using the TPC structure $\mathcal{C}\mathcal{Fuk}(\mathcal{X};\mathcal{P}_2)$ is greater than or equal to the pseudometric defined using $\mathcal{C}\mathcal{Fuk}(\mathcal{X};\mathcal{P}_1)$. By using the functor $\mathcal{F}^{2,1}$ we deduce the opposite inequality, and we conclude that $D^{\mathcal{F}}$ is independent of the perturbation data used to define it.

We begin with the proof of point (i). The inequality relates the spectral distance $\sigma(L,L')$ to the simplest fragmentation metric, $D(-,-) = D^{\{0\}}(-,-)$. The proof is based on a simple consequence of the properties of the Yoneda embedding.

Let $L, L' \in \mathcal{X}$. Without loss of generality, we may assume that $\bar{L}$ and $\bar{L}'$ are Hamiltonian isotopic (otherwise $\sigma(L,L') = +\infty$). Recall also that all Lagrangians in $\mathcal{X}$ are assumed to be graded. Standard Floer theory shows that there is a canonical class $a = o_{L,L'} \in \mathrm{HF}(L,L')$ obtained as the image of the fundamental class $[L] \in H_n(L,\mathbf{k})$ through the PSS morphism $H_*(L,\mathbf{k}) \to \mathrm{HF}(L,L')$. There is a similar class $b = o_{L',L} \in \mathrm{HF}(L',L)$, and it is a simple consequence of the properties of the Yoneda embedding that these two classes have the property that $a * b = [L]$ and $b * a = [L']$. Here $*$ is the product induced in homology by the A_∞-composition μ_2, and $[L]$ and $[L']$ are the respective fundamental classes.

We start with an argument in which we neglect filtration issues. Suppose that there are two classes $a \in \mathrm{HF}(L,L')$ and $b \in \mathrm{HF}(L',L)$ such that $a * b = [e_{L,L}] = [L] \in \mathrm{HF}(L,L)$, where $e_{L,L}$ is the unit in $\mathrm{CF}(L,L)$. For coherence, we work here in homological notation (even if to keep track of signs it would be preferable to use cohomological notation as in [53]), so $e_{L,L}$ corresponds to the maximum of the Morse function $f_{\bar{L}}$ as in Section 3.3.1. This means that there are Floer cycles

$\alpha \in \mathrm{CF}(L, L')$, $\beta \in \mathrm{CF}(L', L)$ such that $\mu_2(\alpha, \beta)$ is homologous to $e_{L,L}$. However, recall that $\mathrm{CF}(L, L)$ is the Morse complex of the function $f_{\bar{L}}$. Thus, for degree reasons we obtain $\mu_2(\alpha, \beta) = e_{L,L}$. As explained in the construction of $\mathcal{F}uk(\mathcal{X}, \mathcal{P})$, $e_{L,L}$ is a strict unit.

For every $L_0, L_1 \in \mathcal{X}$ and any cycle $u \in \mathrm{CF}(L_0, L_1)$, we have a morphism of Yoneda modules $\phi^u : \mathcal{Y}(L_0) \to \mathcal{Y}(L_1)$ defined by

$$\phi^u_k(-, -, \ldots, x) = \mu_{k+1}(-, -, \ldots, x, u).$$

Returning to α and β above, we have that $\phi^{\mu_2(\alpha,\beta)} = \phi^{e_{L,L}} = \mathbb{1}_{\mathcal{Y}(L)}$. The composition $\phi^\beta \circ \phi^\alpha$ is not necessarily equal to the morphism $\phi^{\mu_2(\alpha,\beta)}$, but they are homologous as elements of $\hom_{\mathrm{mod}}(\mathcal{Y}(L), \mathcal{Y}(L))$. (See [53] for the explicit formulas for the composition of pre-module homomorphisms and for the differential μ_1^{mod} on $\hom_{\mathrm{mod}}(-, -)$.) We rewrite here this differential (neglecting signs) for a pre-module morphism $t : \mathcal{M} \to \mathcal{N}$, with components

$$t_k : \mathrm{CF}(L_1, L_2) \otimes \cdots \otimes \mathrm{CF}(L_{k-1}, L_k) \otimes \mathcal{M}(L_k) \to \mathcal{N}(L_1).$$

We have that

$$\begin{aligned}(\mu_1^{\mathrm{mod}} t)_m(-, \ldots, -, x) = &\sum_{r+s=m+1} \mu_r^{\mathcal{N}}(-, -, \ldots, t_s(-, -, \ldots, x)) \\ &+ \sum_{l+j=m+1} t_l(-, -, \ldots, \mu_j^{\mathcal{M}}(-, -, \ldots, x)) \\ &+ \sum_{k+g=m+1} t_k(-, \ldots, \mu_g(-, \ldots, -), -, \ldots, x),\end{aligned}$$

where μ_g is the operation in the A_∞-category, and $\mu_r^{\mathcal{M}}$ and $\mu_j^{\mathcal{N}}$ are the respective module operations. Consider the pre-module morphism $T : \mathcal{Y}(L) \to \mathcal{Y}(L)$ with the kth component given by

$$T_k(-, -, \ldots, -, x) = \mu_{k+2}(-, -, \ldots, -, x, \alpha, \beta).$$

Using the fact that $\mu_2(\alpha, \beta)$ is a strict unit and the A_∞-relations, it is easy to see that $\mu_1^{\mathrm{mod}}(T) = \mathbb{1}_{\mathcal{Y}(L)} - \phi^\beta \circ \phi^\alpha$.

In our case, the A_∞-category $\mathcal{F}uk(\mathcal{X}, \mathcal{P})$ is filtered and, if we assume that $\alpha \in \mathrm{CF}^{\leq r}(L, L')$ and $\beta \in \mathrm{CF}^{\leq s}(L, L')$, our discussion above shows that in the TPC $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ we have maps

$$\Sigma^{r+s}\mathcal{Y}(L) \xrightarrow{\Sigma^r \bar{\phi}^\alpha} \Sigma^s \mathcal{Y}(L') \xrightarrow{\bar{\phi}^\beta} \mathcal{Y}(L)$$

that compose to the "shift" map η_{r+s}. Here $\bar{\phi}^\beta : \Sigma^s \mathcal{Y}(L') \to \mathcal{Y}(L)$ is induced by ϕ^β and $\bar{\phi}^\alpha : \Sigma^r \mathcal{Y}(L) \to \mathcal{Y}(L')$ is induced by ϕ^α. The key ingredient here is that the homotopy T only shifts filtration by at most $r + s$.

Returning to our classes $a = [\alpha]$ and $b = [\beta]$, we now use $b * a = [e_{L',L'}]$ and obtain the existence of a similar diagram:

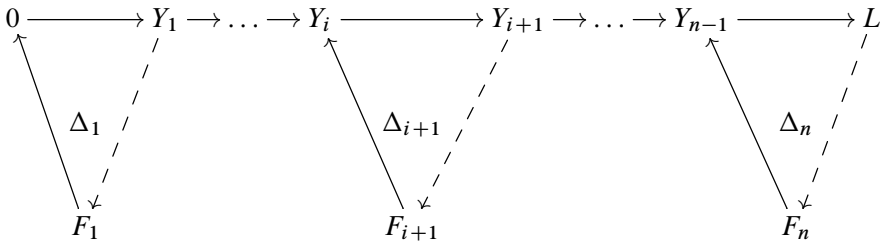
$$\Sigma^{r+s}\mathcal{Y}(L') \xrightarrow{\Sigma^s \bar{\phi}^\beta} \Sigma^r \mathcal{Y}(L) \xrightarrow{\bar{\phi}^\alpha} \mathcal{Y}(L').$$

We immediately deduce from this that $\mathcal{Y}(L)$ and $\mathcal{Y}(L')$ are $(r+s)$-interleaved in the sense of Definition 2.84 in the TPC $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ (one can do better, for instance by using $\max\{r, s\}$, but we will not try to improve the estimates). This means, by Corollary 2.86, that $D(L, L') \leq 4(r+s)$. By reviewing the definition of spectral invariants (2.36), this implies that $D(L, L') \leq 4(\sigma(o_{L,L'}) + \sigma(o_{L',L}))$. The inequality claimed at the first point of the theorem now follows because $\mathrm{HF}(L, L')$ and $\mathrm{HF}(L', L)$ are related by duality and $\sigma(o_{L',L}) = -\sigma(pt_{L,L'})$.

Remark 3.26. (a) The argument at the center of this proof has appeared before in the literature (in less generality, involving only μ_2 and μ_3 operations, instead of the full A_∞-structure) first in [41], as well as in [8]. All these arguments are based on properties of the Yoneda embedding from [53]. These were adjusted to the weakly-filtered setting in [10]. We can achieve a greater degree of generality here because we are able to work with genuinely filtered A_∞-categories which, moreover, are strictly unital.

(b) The fact that the modules $\mathcal{Y}(L)$ and $\mathcal{Y}(L')$ are interleaved by an order controlled by $\sigma(L, L')$ has numerous consequences, because it implies that persistence modules of the form $\mathrm{HF}(N, L)$ and $\mathrm{HF}(N, L')$ are also interleaved by the same order for all $N \in \mathcal{X}$. Viewing this through the prism of the bottleneck metric we easily deduce relations between various measurements of interest, such as boundary depth, spectral range, and so forth, associated to the complexes $\mathrm{CF}(N, L)$ and $\mathrm{CF}(N, L')$.

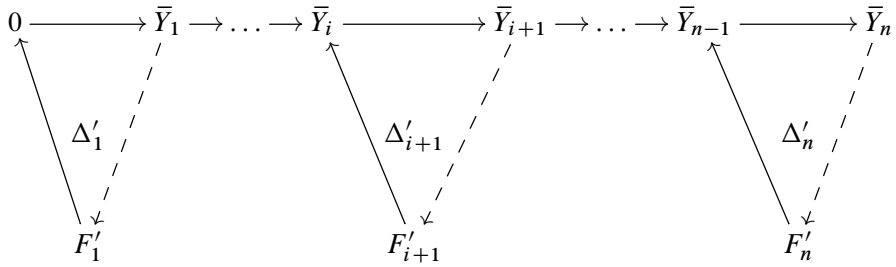
We now turn to the proof of point (ii) in Theorem 3.4. Assume that $D^{\mathcal{F}}(L, L') < r$. Then there exists a sequence of exact triangles in $[\mathcal{C}\mathcal{F}uk(\mathcal{X})]_\infty$ as below:



(3.29)

Here all the terms F_i are in $\mathcal{F}$, except for the jth term which is of the form $T^{-1}\Sigma^\alpha L'$ for some $\alpha \in \mathbb{R}$ (we write here N instead of $\mathcal{Y}(N)$ for the elements of $\mathcal{X}$); see (2.40). Moreover, by assumption we have $\sum_i \bar{w}(\Delta_i) \leq r$. We now appeal to the argument from the proof of Corollary 2.88 to associate to the sequence of exact triangles above

another sequence of triangles:



(3.30)

this time all exact in $[\mathcal{C}\mathcal{F}uk(\mathcal{X})]_0$, together with a $2r$-isomorphism $L \to \bar{Y}_n$. Here the F'_i's are shifts of the previous corresponding objects from (3.29). Given that $\mathcal{F}$ is closed under shifts, we have $F'_i \in \mathcal{F}$ except for the jth term, which is still a shift of $T^{-1}L'$. By possibly shifting $\bar{Y}_n$ down, we have a $4r$-isomorphism $\phi : \bar{Y}_n \to L$.

Each exact triangle Δ_i is the image of a cone-type triangle in the category of filtered modules:

$$F'_{i+1} \xrightarrow{\varphi_i} \bar{Y}_i \to \mathrm{Cone}(\varphi_i) = \bar{Y}_{i+1},$$

and there is a module map $\bar{\phi} : \bar{Y}_n \to L$ such that $K = \mathrm{Cone}(\bar{\phi})$ is $4r$-acyclic. In particular, $K(L)$ is a $4r$-acyclic filtered chain complex: the identity of this chain complex is homotopic to zero through a homotopy h of shift at most $4r$. The complex $K(L)$, viewed as a vector space, is the following sum:

$$\begin{aligned} K(L) &= \mathrm{CF}(L,L) \oplus \bigoplus_{i=1}^{n} \mathrm{CF}(L, F'_i)[2] \\ &= \mathrm{CF}(L,L) \oplus \bigoplus_{i=1}^{j-1} \mathrm{CF}(L, F'_i)[2] \oplus \mathrm{CF}(L,L')[1] \oplus \bigoplus_{i=j+1}^{n} \mathrm{CF}(L, F'_i)[2]. \end{aligned} \tag{3.31}$$

Denote the maximum of the Morse function $f_{\bar{L}} : \bar{L} \to \mathbb{R}$ by m. We now make the following two assumptions:

(i) There exists an embedding of a standard symplectic ball $e' : B(s) \to X$ such that $e'(0) = m$, $(e')^{-1}(L) = \mathbb{R}B(s)$, and $e'(B(s)) \cap (L' \cup \bigcup_{F \in \mathcal{F}} F) = \emptyset$.

(ii) The almost complex structures on X that are part of the perturbation data $\mathcal{P}$ pull back to the standard almost complex structure on $B(s)$.

These assumptions are not restrictive because, as shown at the beginning of the proof, the pseudometric $D^{\mathcal{F}}$ is independent of the perturbation data used in its definition. Thus we can pick perturbation data $\mathcal{P}$ adapted to the embedding e', in the sense that the two points (i) and (ii) above are satisfied, and the sequences of exact triangles (3.29) and (3.30) still exist in this case, with the claimed properties.

Returning to the complex $K(L)$, denote by d_K its differential and consider the equation

$$d_K h(m) + h d_K(m) = m.$$

Now $\mathrm{CF}(L, L)$ is a subcomplex of $K(L)$, and thus $d_K(m) = 0$. Hence $d_K h(m) = m$, so m is a boundary in the complex $(K(L), d_K)$.

On the other hand, since K is an iterated cone of A_∞-Yoneda modules, the differential d_K has a particular shape with respect to the splitting from (3.31). This has been worked out in detail in [10, Section 2.6] (in particular, see Theorem 2.14 in that paper), and the relevant ingredients are as follows. The differential d_K can be described by a matrix $(a_{i,j})_{0 \le i,j \le n}$, where:

(1) $a_{i,j} = 0$ for $i > j$ (i.e., the matrix $(a_{i,j})$ is upper triangular);

(2) $a_{i,j} : \mathrm{CF}(L, F'_j) \to \mathrm{CF}(L, F'_i)$ for $1 \le i \le j \le n$;

(3) $a_{0,j} : \mathrm{CF}(L, F'_j) \to \mathrm{CF}(L, L)$ for $j \ge 1$;

(4) $a_{0,0} : \mathrm{CF}(L, L) \to \mathrm{CF}(L, L)$ is the Floer differential on $\mathrm{CF}(L, L)$.

Here we have omitted reference to the grading on the F'_j's and L. Moreover, for $j \ge 1$, the maps $a_{0,j}$ can be written as follows. There exist Floer chains $c_{q,p} \in \mathrm{CF}(F'_q, F'_p)$ for every $q > p > 0$ and $c_{q,0} \in \mathrm{CF}(F'_q, L)$ for all $q \ge 1$, all at action levels ≤ 0, such that

$$a_{0,j}(-) = \sum_{2 \le d, \underline{k}} \mu_d(-, c_{k_d, k_{d-1}}, \ldots, c_{k_2,0}), \tag{3.32}$$

where $\underline{k} = (k_2, \ldots, k_d)$ runs over all partitions $0 < k_2 < \cdots < k_{d-1} < k_d = j$ and μ_d is the dth-order operation in the Fukaya category $\mathcal{F}uk(\mathcal{X})$.

Since m is the maximum of the Morse function, and m is a boundary in the complex $(K(L), d_K)$, it follows from (3.32) that there is a J-holomorphic polygon with one edge on L and the others on the F'_i's and possibly L' (recall that $F'_j = T^{-1}L'$), that goes through m. (See [10, Section 5.1], and in particular pages 91–92 in that paper, for a detailed proof of a very similar statement.) The area of each such polygon is at least $\frac{1}{2}\pi s^2$. This means that the chain homotopy h increases filtration by at least $\frac{1}{2}\pi s^2$ and thus

$$4r \ge \frac{\pi s^2}{2},$$

which shows the claim.

We now proceed to the proof of point (iii) of Theorem 3.4. We will again make use of the fact that the pseudometric $D^{\mathcal{F}}$ is independent of the perturbation data $\mathcal{P}$.

Let $\delta = \delta^{\cap}(N, L'; \mathcal{F})$. For each intersection point $x \in N \cap L'$, fix standard ball embeddings $e_x : B(s) \to X$ with $e^{-1}(N) = \mathbb{R}B(s)$, $e^{-1}(L') = i\mathbb{R}B(s)$, $e(0) = x$, such that all these embeddings are disjoint from the family $\mathcal{F}$ and additionally $\pi s^2 = \delta - \epsilon$ for a small ϵ. We may assume that the almost complex structures that are part of the perturbation data $\mathcal{P}$ pull back to the standard almost complex structure through

the embeddings e_x. This implies that if a Floer-type strip or polygon has an input or an output at a point of $N \cap L'$ and has boundary on N, L', and any other elements of the family $\mathcal{F}$, then its energy is at least $\delta' = \frac{\pi s^2}{4} - \epsilon' = \frac{\delta - \epsilon}{4} - \epsilon'$. Here the small ϵ' has to do with making the Hamiltonian (or 1-forms) part of the perturbation data small enough.

Assume now, as in the statement, that the condition $D^{\mathcal{F}}(L, L') < r < \frac{\delta}{16}$ holds. Then there exist a sequence of exact triangles in $\mathcal{C}\mathcal{F}uk(\mathcal{X})_0$ as in (3.30) and a $2r$-isomorphism $\psi : L \to \bar{Y}_n$. By Lemma 2.85 this means that $d_{\text{int}}(L, Y_n) \leq 2r$. In particular, there exist maps $u : \Sigma^r \bar{Y}_n \to L$, $v : \Sigma^r L \to \bar{Y}_n$ such that $v \circ \Sigma^r u = \eta_{2r}$. Therefore, given that $N \in \mathcal{X}$, we obtain maps of filtered complexes

$$\bar{Y}_n(N) \xrightarrow{u'} \mathrm{CF}(N, L) \xrightarrow{v'} \bar{Y}_n(N),$$

each of shift at most $2r$, and whose composition is chain homotopic to the identity through a homotopy h that shifts filtration by at most $4r$.

Consider the differential of the complex $\bar{Y}_n(N)$. As a vector space, the complex $\bar{Y}_n(N)$ is a sum of the form $\mathrm{CF}(N, F_1') \oplus \cdots \oplus \mathrm{CF}(N, L') \oplus \mathrm{CF}(N, F_n')$, and the differential is represented by clustered polygons with boundaries on N, L', and the elements of the family $\mathcal{F}$.

Consider the composition $\Psi = p \circ v' \circ u' \circ i$, where $i : \mathrm{CF}(N, L') \to \bar{Y}_n(N)$ is the inclusion and $p : \bar{Y}_n(N) \to \mathrm{CF}(N, L')$ is the projection – in both cases as vector spaces. We claim that Ψ is injective. This would imply that $\dim_{\mathbf{k}} \mathrm{CF}(N, L') \leq \dim_{\mathbf{k}} \mathrm{CF}(N, L)$ and prove the statement in point (iii) of Theorem 3.4.

To show the injectivity of Ψ, we inspect the formula $dh + hd = \mathbb{1} - v' \circ u'$, and recall that the differential d, when restricted to $\mathrm{CF}(N, L')$, drops the filtration by at least $\delta' = \frac{\delta - \epsilon}{4} - \epsilon'$ while h raises filtration by at most $4r < \delta'$ (when ϵ and ϵ' are small enough). In other words, $p(hd(x))$ for each element $x \in \mathrm{CF}(N, L')$ is of strictly lower filtration than x. Similarly, $p(dh(x))$ is also of lower filtration than x. As a result, we deduce that Ψ can be written as $\mathbb{1}$ plus a map that strictly lowers the filtration level. It follows that Ψ is injective.

Finally, we discuss the last point in Theorem 3.4. We assume that $\mathcal{F}$ generates the usual (i.e., without persistence structure) derived Fukaya category $D\mathcal{F}uk(X)$ and we want to show that in this case the pseudometric $D^{\mathcal{F}}$ is finite. From Theorem 3.1 (iii) we deduce that $\mathcal{F}$ generates $\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty$. This implies that any object A in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ is r-isomorphic, for some r, to an object that can be written as an iterated cone with triangles of weight 0, of the form in (3.30). This means that $D^{\mathcal{F}}(A, 0) < \infty$ and implies the last claim in the statement of the theorem. ∎

3.4.2 Pseudometrics on $\mathcal{Lag}(X)$ and proof of Corollary 3.7

The purpose of this subsection is to use the results from Section 3.1 to construct the family of fragmentation metrics on the space $\mathcal{Lag}(X)$ of all closed exact, graded Lagrangians in (X, ω) and prove Corollary 3.7.

We assume the setting in Corollary 3.7, in particular that rank $\mathcal{Fuk}(X, \omega) < \infty$. Thus $D\mathcal{Fuk}(X)$ admits a finite set of triangular generators, in the sense that there is a family of triangular generators $\mathcal{F}$ such that the corresponding family $\overline{\mathcal{F}} \subset \mathcal{Lag}(X)$, obtained by forgetting the grading and the choices of primitives, is finite. We now fix such a family $\mathcal{F}$ of generators and assume that $\mathcal{F}$ is invariant under shifts and translations, and that the Lagrangians in $\overline{\mathcal{F}}$ are in general position (every two Lagrangians intersect transversely and there are no triple intersection points).

The proof of Corollary 3.7 is a consequence of Theorem 3.4 together with the invariance properties of the Fukaya TPC constructed earlier in this section, and is contained in the subsections below.

3.4.2.1 The Fukaya TPC revisited. Pick a shift- and translation-invariant family $\mathcal{X} \subset \mathcal{Lag}(X)'$ that contains $\mathcal{F}$. As before, we denote by $\overline{\mathcal{X}} \subset \mathcal{Lag}(X)$ the corresponding family of Lagrangians after forgetting the choices of primitives and grading. We assume that $\overline{\mathcal{X}}$ is finite and that its elements are in general position.

We recall the filtered Fukaya category $\mathcal{Fuk}(\mathcal{X}; \mathcal{P})$, where $\mathcal{P}$ is a choice of perturbation data, as in Theorem 3.12; see also Section 3.4.1.1. We already know that any two such categories, defined using two admissible perturbation data $\mathcal{P}$ and $\mathcal{P}'$, are filtered quasi-equivalent, in the sense that there are filtered A_∞-functors $\mathcal{Fuk}(\mathcal{X}; \mathcal{P}) \to \mathcal{Fuk}(\mathcal{X}; \mathcal{P}')$ that are the identity on objects and induce a (filtered) equivalence of the homological persistence categories.

As discussed in Section 3.2.1, there are two TPCs that one can associate to $\mathcal{Fuk}(\mathcal{X}; \mathcal{P})$. The first is $\mathcal{C}\mathcal{Fuk}(\mathcal{X})$ as in Theorem 3.1. This is obtained by considering the filtered modules $\mathrm{Fmod}(\mathcal{Fuk}(\mathcal{X}; \mathcal{P}))$ over $\mathcal{Fuk}(\mathcal{X}; \mathcal{P})$. The category $\mathcal{C}\mathcal{Fuk}(\mathcal{X})$ is the homological category

$$\mathcal{C}\mathcal{Fuk}(\mathcal{X}) = H^0\big[\mathcal{Fuk}(\mathcal{X}; \mathcal{P})^\nabla\big],$$

where $\mathcal{Fuk}(\mathcal{X}; \mathcal{P})^\nabla$ is the smallest triangulated (with respect to weight-0 triangles) full subcategory of $\mathrm{Fmod}(\mathcal{Fuk}(\mathcal{X}; \mathcal{P}))$ that contains the Yoneda modules $\mathcal{Y}(L)$ with $L \in \mathcal{X}$ and is closed under r-isomorphisms for all r, in the sense that if $j : M \to M'$ is an r-isomorphism of modules and $M \in \mathcal{Fuk}(\mathcal{X}; \mathcal{P})^\nabla$, then $M' \in \mathcal{Fuk}(\mathcal{X}; \mathcal{P})^\nabla$. Possibly more concretely, each object in this category is r-isomorphic, for some r, to a weight-0 iterated cone of Yoneda modules. Two such triangulated persistence categories, defined for different choices of perturbation data, are TPC equivalent (see also Remark 3.2), and thus we drop the reference to the perturbation data from the notation.

The second type of TPC will be denoted by $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$ and is defined by

$$\mathcal{C}'\mathcal{F}uk(\mathcal{X}) = H^0[\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))], \tag{3.33}$$

where $\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$ is the category of filtered twisted complexes constructed from $\mathcal{F}uk(\mathcal{X};\mathcal{P})$; see Section 3.2.1. In this case too we can drop the reference to the choices of perturbation data, as any two such choices produce equivalent TPCs (again, these equivalences are not entirely canonical, but this will not have any impact on our further arguments).

There are filtered functors

$$\Theta : \mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P})) \to \mathrm{Fmod}[\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))] \to \mathrm{Fmod}(\mathcal{F}uk(\mathcal{X};\mathcal{P})),$$

where the first arrow is the Yoneda embedding and the second is the pullback over the natural inclusion $\mathcal{F}uk(\mathcal{X};\mathcal{P}) \to \mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$. The composition Θ is a homologically full and faithful embedding, and it induces a full and faithful embedding of TPCs. The image of Θ lands inside $\mathcal{F}uk(\mathcal{X})^\nabla$ (actually inside the category denoted by $\mathcal{F}uk(\mathcal{X},\mathcal{P})^\#$ in Section 3.2.1), and thus we have an inclusion of TPCs:

$$\bar{\Theta} : \mathcal{C}'\mathcal{F}uk(\mathcal{X}) \hookrightarrow \mathcal{C}\mathcal{F}uk(\mathcal{X}).$$

By contrast with the unfiltered case, $\bar{\Theta}$ is not an equivalence of TPCs because, for it to be an equivalence of TPCs, each object in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ needs to be 0-isomorphic to some object in $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$ (see Definition 2.25) and, a priori, this might not happen (each object in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ is r-isomorphic to some object in $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$, but possibly for $r > 0$).

3.4.2.2 Fragmentation metrics as in Theorem 3.4. We now focus on the pseudometric $D^{\mathcal{F}}$ as in Theorem 3.4. This is constructed by the general procedure described in the algebraic part of this memoir, by using the persistence triangular weight on $\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty$. Thus, this is a shift-invariant fragmentation pseudometric of the type $\widehat{d^{\mathcal{F}}}(-,-)$, as described in Section 2.4.3.1, restricted to $\mathcal{X}$. To emphasize the relation of this pseudometric to the set $\mathcal{X}$ we will denote it by $D^{\mathcal{F}}_{\mathcal{X}}$. Because the pseudometric is shift invariant, it descends to $\bar{\mathcal{X}}$. As shown in Theorem 3.4, this pseudometric is independent of the choice of perturbation data $\mathcal{P}$.

Another way to construct a fragmentation-type pseudometric on $\bar{\mathcal{X}}$ is to apply the exact same construction to the triangulated persistence category $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$ from (3.33). The resulting pseudometric, again defined on $\mathcal{X}$, will be denoted by $\bar{D}^{\mathcal{F}}_{\mathcal{X}}$. It is again independent of the perturbation data used in its definition. If one is interested only in comparing objects in $\mathcal{X}$, this pseudometric is much more approachable from a computational point of view, because $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$ has fewer objects than $\mathcal{C}\mathcal{F}uk(\mathcal{X})$.

There is a simple relation between the two pseudometrics discussed above.

Lemma 3.27. *For any $L, L' \in \bar{\mathcal{X}}$, the pseudometrics $\bar{D}^{\mathcal{F}}_{\mathcal{X}}$ and $D^{\mathcal{F}}_{\mathcal{X}}$ satisfy the inequalities*

$$\frac{1}{4}\,\bar{D}^{\mathcal{F}}_{\mathcal{X}}(L, L') \leq D^{\mathcal{F}}_{\mathcal{X}}(L, L') \leq \bar{D}^{\mathcal{F}}_{\mathcal{X}}(L, L').$$

Proof. Indeed, the inequality on the right is obvious in view of the TPC inclusion $\bar{\Theta}$. The inequality on the left follows from the argument in Lemma 2.87. In fact, this argument shows that if we have a cone decomposition of $\mathcal{Y}(L)$ of total weight r in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$, with linearization consisting of Yoneda modules of elements in $\mathcal{F}$ together with one term that is a shift/translate of $\mathcal{Y}(L')$, then there is another cone decomposition of $\mathcal{Y}(L)$ with a linearization of the same type such that all triangles are of weight 0 except the last triangle, which is of weight at most $4r$. But this means that all these triangles can be assumed to belong to $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$. This is clear for the weight-0 triangles because the image of Θ contains the Yoneda modules $\mathcal{Y}(L)$, $L \in \mathcal{X}$, and is closed with respect to taking cones over filtration-preserving maps. The last triangle corresponds to a strict exact triangle of weight $4r$ with the $4r$-isomorphism on the third term of the form $\phi : M_{n-1} \to \mathcal{Y}(L)$, where M_{n-1} is an object of $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$. Because $\bar{\Theta}$ is full and faithful, we have that $\phi \in \hom_{\mathcal{C}'\mathcal{F}uk(\mathcal{X})}$. From this argument, it is easy to deduce the inequality on the left in the statement of the lemma, which concludes the proof. ∎

3.4.2.3 Changing the set $\mathcal{X}$. Assume that $\mathcal{X}'$ is another family of marked Lagrangians which is invariant under shifts and translations and such that $\mathcal{X} \subset \mathcal{X}'$. Additionally, we assume that the family $\bar{\mathcal{X}}'$ is in general position.

Lemma 3.28. *Under the assumptions above, and for any two $L, L' \in \bar{\mathcal{X}}$, we have*

$$\bar{D}^{\mathcal{F}}_{\mathcal{X}}(L, L') \geq \bar{D}^{\mathcal{F}}_{\mathcal{X}'}(L, L') \quad \text{and} \quad D^{\mathcal{F}}_{\mathcal{X}}(L, L') \leq D^{\mathcal{F}}_{\mathcal{X}'}(L, L').$$

Proof. We consider the category $\mathcal{F}uk(\mathcal{X}', \mathcal{P}')$ and we notice that the restriction of $\mathcal{P}'$ to $\mathcal{X}$ provides an allowable choice of perturbation data; see Theorem 3.12. Thus, there is an A_∞ filtered inclusion

$$\mathcal{F}uk(\mathcal{X}, \mathcal{P}') \to \mathcal{F}uk(\mathcal{X}', \mathcal{P}).$$

This inclusion induces a pullback TPC functor

$$qJ^* : \mathcal{C}\mathcal{F}uk(\mathcal{X}') \to \mathcal{C}\mathcal{F}uk(\mathcal{X}).$$

This is well defined because $\mathcal{F} \subset \mathcal{X} \subset \mathcal{X}'$ is a system of triangular generators for $D\mathcal{F}uk(X)$, which means that, in particular, any Yoneda module $\mathcal{Y}(L)$, $L \in \mathcal{X}'$, is r-isomorphic (in $\mathcal{C}\mathcal{F}uk(\mathcal{X}')$), for some $r \geq 0$, to a 0-weight iterated cone of Yoneda modules of elements from $\mathcal{F}$. This means that the pullback of $\mathcal{Y}(L)$ to $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X}, \mathcal{P}'))$ is an object of $\mathcal{C}\mathcal{F}uk(\mathcal{X})$.

The same A_∞ inclusion also induces a pushforward TPC functor

$$J_* : \mathcal{C}'\mathcal{F}uk(\mathcal{X}) \to \mathcal{C}'\mathcal{F}uk(\mathcal{X}'),$$

which is induced by the natural inclusion of twisted complexes.

As discussed before, our invariance statements imply that the pseudometrics $\bar{D}^{\mathcal{F}}_{\mathcal{X}}$ and $D^{\mathcal{F}}_{\mathcal{X}}$ do not depend on the choice of perturbation data. As a result, the fact that J_* is a TPC functor implies the first inequality in the lemma, and the fact that J^* is a TPC functor implies the second inequality. ∎

3.4.2.4 The pseudometric $\mathcal{D}^{\mathcal{F}}$ from Corollary 3.7. The construction of $\mathcal{D}^{\mathcal{F}}$ proceeds in two steps.

The first step is to consider again a family $\mathcal{X}$ as in the subsections above, together with two elements $L, L' \in \mathcal{L}ag(X)$. We define

$$\mathcal{D}^{\mathcal{F}}_{\mathcal{X}}(L, L') = \limsup_{\epsilon \to 0} D^{\mathcal{F}}_{\mathcal{X} \cup \{L_\epsilon, L'_\epsilon\}}(L_\epsilon, L'_\epsilon).$$

Here L_ϵ and L'_ϵ are Hamiltonian deformations of, respectively, L and L' through Hamiltonians of Hofer norm at most $\epsilon \geq 0$, such that the family $\mathcal{X} \cup \{L_\epsilon, L'_\epsilon\}$ is allowable for the definition of the Fukaya categories $\mathcal{F}uk(\mathcal{X} \cup \{L_\epsilon, L'_\epsilon\}; \mathcal{P})$. The lim sup is taken over all possible choices of such perturbations as ϵ goes to 0.

The second step is to put

$$\mathcal{D}^{\mathcal{F}}(L, L') = \sup_{\mathcal{X}} \mathcal{D}^{\mathcal{F}}_{\mathcal{X}}(L, L').$$

It is clear that $\mathcal{D}^{\mathcal{F}}$ is symmetric. We will see below, in Lemma 3.30, that $\mathcal{D}^{\mathcal{F}}$ is finite. However, before we get to that, we have the following result.

Lemma 3.29. *With the definition above, $\mathcal{D}^{\mathcal{F}}$ satisfies the triangle inequality.*

Proof. We fix three Lagrangians L, L', L'' in $\mathcal{L}ag(X)$. Fix also a family $\mathcal{X}$ as above. Consider L_n, L''_n in $\mathcal{L}ag(X)$ such that the family $\mathcal{X} \cup \bigcup_{n,m}\{L_n, L''_m\}$ is in general position in our usual sense: any couple of Lagrangians in the family intersect transversely and there are no triple intersection points (this choice is possible as there are only countable many transversality-type constraints). We also assume:

- $d_H(L, L_n) \leq \frac{1}{n}$ and $d_H(L'', L''_n) \leq \frac{1}{n}$, where $d_H(-,-)$ is the Hofer distance;
- $\lim_{n\to\infty} D^{\mathcal{F}}_{\mathcal{X} \cup \{L_n, L''_n\}}(L_n, L''_n) = \mathcal{D}^{\mathcal{F}}_{\mathcal{X}}(L, L'')$.

The lemma would follow if we prove that, for every $\delta > 0$,

$$\mathcal{D}^{\mathcal{F}}_{\mathcal{X}}(L, L'') \leq \mathcal{D}^{\mathcal{F}}(L, L') + \mathcal{D}^{\mathcal{F}}(L', L'') + 4\delta. \tag{3.34}$$

To show (3.34), we pick a sequence of Lagrangians L'_k such that the family $\mathcal{X} \cup \bigcup_{n,k,m}\{L_n, L'_k, L''_m\}$ is in general position and $d_H(L', L'_n) \leq \frac{1}{n}$. For any m, n, k, we have the inequalities

$$\begin{aligned} D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L''_m\}}(L_n, L''_m) &\leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_m\}}(L_n, L''_m) \\ &\leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_m\}}(L_n, L'_k) + D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_m\}}(L'_k, L''_m). \end{aligned}$$

The first inequality comes from Lemma 3.28 and the second is the triangle inequality for the fragmentation pseudometric $D^{\mathcal{F}}_{(-)}$. We will estimate separately the two terms on the right-hand side of this inequality.

Fix a natural number m_0. We can find $N_{m_0} \geq m_0$ such that for $n, k \geq N_{m_0}$ we have

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0}\}}(L_n, L'_k) \leq \mathcal{D}^{\mathcal{F}}_{\mathcal{X}\cup\{L''_{m_0}\}}(L, L') + \delta \leq \mathcal{D}^{\mathcal{F}}(L, L') + \delta.$$

Thus, for $n, k \geq N_{m_0}$ we have

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L''_{m_0}\}}(L_n, L''_{m_0}) \leq \mathcal{D}^{\mathcal{F}}(L, L') + \delta + D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0}\}}(L'_k, L''_{m_0}),$$

and it remains to estimate the rightmost term. Using Lemma 3.28 again and the triangle inequality, we have

$$\begin{aligned} &D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0}\}}(L'_k, L''_{m_0}) \\ &\quad\leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L'_k, L''_{m_0}) \\ &\quad\leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L'_k, L''_m) + D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L''_m, L''_{m_0}). \end{aligned}$$

All our fragmentation pseudometrics are bounded above by the Hofer norm and thus we have $D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L''_m, L''_{m_0}) \leq \frac{2}{m_0}$ as soon as $m \geq m_0$. We now consider $n \geq N_{m_0}$, and take k, m sufficiently large such that we have

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L'_k, L''_m) \leq \mathcal{D}^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L''_{m_0}\}}(L', L'') + \delta \leq \mathcal{D}^{\mathcal{F}}(L', L'') + \delta.$$

Putting things together, for our fixed (arbitrary) m_0 and any $n \geq N_{m_0}$ we have

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L''_{m_0}\}}(L_n, L''_{m_0}) \leq \mathcal{D}^{\mathcal{F}}(L, L') + \delta + \mathcal{D}^{\mathcal{F}}(L', L'') + \delta + \frac{2}{m_0}. \tag{3.35}$$

An important remark is that inequality (3.35) applies to any fixed m_0 and any set $\mathcal{X}'$ that contains $\mathcal{X}$ (and is such that the family $\mathcal{X}' \cup \bigcup_{n,k,m}\{L_n, L'_k, L''_m\}$ is in general position), since the argument above applies verbatim to this situation. In this case the number N_{m_0} depends on m_0 but also, implicitly, on $\mathcal{X}'$. To make this dependence explicit we will denote it by $N_{m_0,\mathcal{X}'}$.

For any n, we have the triangle inequality:

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_{m_0},L''_{m_0}) \\ \leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_n,L''_{m_0}) + D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_{m_0},L_n).$$

Assuming $n \geq m_0$, the second term on the right-hand side is bounded above by $\frac{2}{m_0}$ and thus, using Lemma 3.28, we deduce

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L''_{m_0}\}}(L_{m_0},L''_{m_0}) \leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_{m_0},L''_{m_0}) \\ \leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_n,L''_{m_0}) + \frac{2}{m_0}.$$

We now apply (3.35) to $\mathcal{X}' = \mathcal{X} \cup \{L_{m_0}\}$. We deduce that, for $n \geq N_{m_0,\mathcal{X}'}$,

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_n,L''_{m_0}) \leq \mathcal{D}^{\mathcal{F}}(L,L') + \mathcal{D}^{\mathcal{F}}(L',L'') + 2\delta + \frac{2}{m_0}.$$

The last inequality implies

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L''_{m_0}\}}(L_{m_0},L''_{m_0}) \leq \mathcal{D}^{\mathcal{F}}(L,L') + \mathcal{D}^{\mathcal{F}}(L',L'') + 2\delta + \frac{4}{m_0}.$$

As this is true for an arbitrary choice of m_0, we deduce inequality (3.34). This concludes the proof. ∎

3.4.2.5 Properties of $\mathcal{D}^{\mathcal{F}}$. We know from Section 3.4.2.4 that $\mathcal{D}^{\mathcal{F}}$ is a pseudometric. In this subsection we will show that $\mathcal{D}^{\mathcal{F}}$ satisfies the other properties claimed in Corollary 3.7. The properties (i), (ii), and (iii), are in fact immediate consequences of the properties of the pseudometrics $D^{\mathcal{F}}_{\mathcal{X}}$ that appear in Theorem 3.4. Indeed, the estimates in this theorem do not depend on the set $\mathcal{X}$, and this easily implies the corresponding properties for $\mathcal{D}^{\mathcal{F}}$. A more delicate property is the following one.

Lemma 3.30. *The pseudometric $\mathcal{D}^{\mathcal{F}}$ is finite.*

Proof. This property follows by applying repeatedly Lemmas 3.27 and 3.28. Fix $L, L' \in \mathcal{L}ag(X)$, a family $\mathcal{X}$ as before, and perturbations L_ϵ, L'_ϵ such that $\mathcal{X} \cup \{L_\epsilon, L'_\epsilon\}$ is in general position and L_ϵ, L'_ϵ are ϵ-close to L and L', respectively, in the Hofer norm. We have

$$\bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L'_\epsilon\}}(L_\epsilon,L'_\epsilon) \geq \bar{D}^{\mathcal{F}}_{\mathcal{X}\cup\{L_\epsilon,L'_\epsilon\}}(L_\epsilon,L'_\epsilon) \geq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_\epsilon,L'_\epsilon\}}(L_\epsilon,L'_\epsilon).$$

Thus the argument reduces to showing that $\bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L'_\epsilon\}}(L_\epsilon,L'_\epsilon)$ has a uniform upper bound when $\epsilon \to 0$. It is immediate to see that it is enough to find such a bound for

$$\bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon\}}(L_\epsilon,0).$$

We fix some ϵ_0 and write, for $\epsilon \leq \epsilon_0$,

$$\begin{aligned}\frac{1}{4}\bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon\}}(L_\epsilon, 0) &\leq D^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon\}}(L_\epsilon, 0) \leq D^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon, L_{\epsilon_0}\}}(L_\epsilon, 0)\\ &\leq \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon, L_{\epsilon_0}\}}(L_\epsilon, 0)\\ &\leq \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon, L_{\epsilon_0}\}}(L_{\epsilon_0}, 0) + \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon, L_{\epsilon_0}\}}(L_{\epsilon_0}, L_\epsilon)\\ &\leq \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon, L_{\epsilon_0}\}}(L_{\epsilon_0}, 0) + 2\epsilon_0\\ &\leq \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_{\epsilon_0}\}}(L_{\epsilon_0}, 0) + 2\epsilon_0.\end{aligned}$$

The first three inequalities come from Lemmas 3.27 and 3.28, and, the fourth from the triangle inequality. The next inequality is implied by the upper bound given by the Hofer norm. Finally, the last inequality comes from Lemma 3.28. ∎

The following result concludes the proof of Corollary 3.7.

Lemma 3.31. *If $\mathcal{F}'$ is a generic perturbation of $\mathcal{F}$ (in the sense that each element of $\mathcal{F}'$ is a small Hamiltonian perturbation of a corresponding element in $\mathcal{F}$ and the union $\mathcal{F} \cup \mathcal{F}'$ is in general position), then*

$$\mathcal{D}^{\mathcal{F},\mathcal{F}'} = \max\{\mathcal{D}^{\mathcal{F}}, \mathcal{D}^{\mathcal{F}'}\}$$

is non-degenerate.

Proof. The statement follows from point (ii) of Corollary 3.7, which has already been shown. From this point we deduce that $\mathcal{D}^{\mathcal{F},\mathcal{F}'}(L, L') = 0$ implies

$$\delta(L; L' \cup \textstyle\bigcup_{F\in\mathcal{F}} F) = 0 = \delta(L; L' \cup \bigcup_{F'\in\mathcal{F}'} F').$$

The definition of $\delta(-;-)$ is in (3.1). The first of the last two equalities means that there is no standard symplectic ball of any positive radius with its real part on L and that is disjoint from $L' \cup \bigcup_{F\in\mathcal{F}} F$. It follows that $L \subset L' \cup \bigcup_{F\in\mathcal{F}} F$. Of course the metric is symmetric, so we also have that $L' \subset L \cup \bigcup_{F\in\mathcal{F}} F$. Given that the same relations are valid for $\mathcal{F}'$, and that $\mathcal{F}'$ and $\mathcal{F}$ are in general position, we deduce that $L = L'$. ∎

3.5 The geometry behind TPCs

This section illustrates geometrically some of the TPC machinery. It contains two subsections. In Section 3.5.1 we explain how the theory of Lagrangian cobordism provides a concrete representation for the algebraic structures that are formalized in the language of TPCs. Some of the material presented in Section 3.5.1 is based on the theory developed in [7, 10], which is technically involved. Below, however, we have

tried to avoid technicalities as much as possible in order to focus on the geometric ideas, at the expense of skipping some details. Most of these can be found in the above references.

In Section 3.5.2 we work out some examples where the estimates in Theorem 3.4 can be made concrete.

3.5.1 Lagrangian cobordism and weighted triangles

The theory of Lagrangian cobordism exhibits in a geometric way several key notions that are fundamental for the algebraic theory of TPCs. The purpose of this section is to provide some geometric interpretations of these notions – in particular, of weighted exact triangles – by using certain natural symplectic measurements associated to Lagrangian cobordisms. The geometric weights coming from geometry (such as those reflecting the shadows of cobordisms) are bigger than the algebraic weights discussed before in this memoir. The difficulty with using them in practice is that they depend on constructing specific cobordisms.

3.5.1.1 Background on Lagrangian cobordism. Let $(X, \omega = d\lambda)$ be a Liouville manifold, endowed with a given Liouville form λ. We endow $\mathbb{R}^2$ with the Liouville form $\lambda_{\mathbb{R}^2} = xdy$ and its associated standard symplectic structure $\omega_{\mathbb{R}^2} = d\lambda_{\mathbb{R}^2} = dx \wedge dy$. Let $\widetilde{X} := \mathbb{R}^2 \times X$, endowed with the Liouville form $\widetilde{\lambda} = \lambda_{\mathbb{R}^2} \oplus \lambda$ and the symplectic structure $\widetilde{\omega} = d(\widetilde{\lambda}) = \omega_{\mathbb{R}^2} \oplus \omega$. We denote by $\pi : \mathbb{R}^2 \times X \to X$ the projection.

Below we will assume known the notion of Lagrangian cobordism, as developed in [6, 7]. For simplicity, we will consider only *negative-ended* cobordisms $V \subset \mathbb{R}^2 \times X$, which means that V has only negative ends. Moreover, all the cobordisms considered below will be assumed to be exact with respect to the Liouville form $\widetilde{\lambda}$ and endowed with a given primitive $f_V : V \to \mathbb{R}$ of $\widetilde{\lambda}|_V$. Denote by $L_1, \dots, L_k \subset X$ the Lagrangians corresponding to the ends of V and by $\ell_1, \dots, \ell_k \subset \mathbb{R}^2$ the negative horizontal rays of V, so that V coincides at $-\infty$ with $(\ell_1 \times L_1) \coprod \cdots \coprod (\ell_k \times L_k)$. We remark that we adopt here the conventions from [7] regarding the ends of V, namely we always assume that the jth ray ℓ_j lies on the horizontal line $\{y = j\}$. Also, we allow some of the Lagrangians L_j to be void.

Note that $\lambda_{\mathbb{R}^2}|_{\ell_i} = 0$, hence $f_V|_{\ell_i \times L_i}$ is constant in the ℓ_i direction for all i. Therefore the Lagrangians $L_i \subset X$ are λ-exact and f_V induces well-defined primitives $f_{L_i} : L_i \to \mathbb{R}$ of $\lambda|_{L_i}$ for each i, namely $f_{L_i}(p) := f_V(z_0, p)$ for every $p \in L_i$, where z_0 is any point on ℓ_i.

3.5.1.2 Weakly filtered Fukaya categories and cobordism. As constructed in [8, 10], there is a weakly filtered Fukaya A_∞-category $W\mathcal{F}uk(X)$ of λ-exact Lagrangians whose objects are exact Lagrangians $L \subset X$ endowed with a primitive $f_L : L \to \mathbb{R}$

of $\lambda|_L$. The notation is ad hoc here to distinguish this category from the filtered versions constructed in Section 3.3.

Remark 3.32. The filtered A_∞ categories $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ from Section 3.3 (see Theorem 3.12) are constructed under more restrictive assumptions than $W\mathcal{F}uk(X)$, since they are associated to only a finite number of geometric objects $\overline{\mathcal{X}}$. However, it is easy to see that, by choosing the perturbation data required to define the weakly filtered category $W\mathcal{F}uk(X)$ so that it extends the data $\mathcal{P}$, we have an embedding (of weakly filtered A_∞-categories)

$$\mathcal{F}uk(\mathcal{X};\mathcal{P}) \to W\mathcal{F}uk(X).$$

There is also a weakly filtered Fukaya A_∞-category of cobordisms $W\mathcal{F}uk(\mathbb{R}^2\times X)$ whose objects are negative-ended exact cobordisms $V \subset \mathbb{R}^2 \times X$ endowed with a primitive f_V of $\widetilde{\lambda}|_V$, as above.[1] We also have the dg-categories of weakly filtered A_∞-modules over each of the previous Fukaya categories, which we denote, respectively, by $\mathrm{mod}_{W\mathcal{F}uk(X)}$ and $\mathrm{mod}_{W\mathcal{F}uk(\mathbb{R}^2\times X)}$.

Below we will mostly concentrate on the chain complexes associated to various Lagrangians and modules, ignoring the higher-order A_∞-operations, and these are genuinely filtered. Thus, in this discussion the fact that the above categories are only weakly filtered rather than genuinely filtered will not a play an important role.

Let $\mathcal{Y} : W\mathcal{F}uk(X) \to \mathrm{mod}_{W\mathcal{F}uk(X)}$ be the Yoneda embedding (in the framework of weakly filtered A_∞-categories), and let $W\mathcal{F}uk(X)^\nabla \subset \mathrm{mod}_{W\mathcal{F}uk(X)}$ be the triangulated closure of the image of $\mathcal{Y}$. We denote by $\mathcal{C} = \mathcal{P}H(\mathcal{F}uk(X)^\nabla)$ the persistence homological category associated to $\mathcal{F}uk(X)^\nabla$. This is not a TPC due to the difference between "filtered" and "weakly filtered"; however, with this distinction kept in mind, its properties closely mimic those of a TPC. To understand the difference, while in a TPC the composition of two morphisms f of shift r and g of shift s is a morphism $f \circ g$ of shift $r + s$, in the weakly filtered case the composition $f \circ g$ has shift $r + s + \epsilon^{\mu_2}$, where the error term μ_2 is part of the structural data associated to the *weakly* filtered structure of $W\mathcal{F}uk(X)$.

By a slight abuse of notation we will denote the Yoneda module $\mathcal{Y}(L)$ of a Lagrangian $L \in \mathrm{Obj}(W\mathcal{F}uk(X))$ also by L.

There is also a Yoneda embedding $W\mathcal{F}uk(\mathbb{R}^2 \times X) \to \mathrm{mod}_{W\mathcal{F}uk(\mathbb{R}^2\times X)}$ and we will typically denote the Yoneda modules corresponding to cobordisms by calligraphic letters, e.g., the Yoneda module corresponding to $V \in \mathrm{Obj}(\mathcal{F}uk(\mathbb{R}^2 \times X))$ will be denoted by $\mathcal{V}$.

Under additional assumptions on X, on the Lagrangians taken as the objects of $W\mathcal{F}uk(X)$, and on the Lagrangian cobordisms of $\mathcal{F}uk(\mathbb{R}^2 \times X)$, one can set up

[1] For technical reasons one needs to enlarge the set of objects in $\mathcal{F}uk(\mathbb{R}^2 \times X)$ to contain also objects of the type $\gamma \times L$, where $\gamma \subset \mathbb{R}^2$ is a curve which, outside a compact set, is either vertical or coincides with horizontal ends with y-value being $l \pm \frac{1}{10}$, where $l \in \mathbb{Z}$.

a graded theory, endowing the morphisms in $\mathcal{F}uk(X)$ and $\mathcal{F}uk(\mathbb{R}^2 \times X)$ with a $\mathbb{Z}$-grading and the categories with grading-translation functors. See [53] for the case of $W\mathcal{F}uk(X)$ and [38] for grading in the framework of cobordisms. In what follows we will not work explicitly in a graded setting, but whenever possible we will indicate how grading fits in various constructions.

3.5.1.3 Iterated cones associated to cobordisms. Let $\gamma \subset \mathbb{R}^2$ be an oriented[2] plane curve that is the image of a proper embedding of $\mathbb{R}$ into $\mathbb{R}^2$. Viewing $\gamma \subset \mathbb{R}^2$ as an exact Lagrangian we fix a primitive f_γ of $\lambda_{\mathbb{R}^2}|_\gamma$. Given an exact Lagrangian $L \subset X$, consider the exact Lagrangian $\gamma \times L \subset \mathbb{R}^2 \times X$ and take $f_{\gamma,L} : \gamma \times L \to \mathbb{R}$, $f_{\gamma,L}(z, p) = f_\gamma(z) + f_L(p)$ for the primitive of $\widetilde{\lambda}|_{\gamma\times L}$. From now on we will make the following additional assumptions on γ. The ends of γ will be assumed to coincide with a pair of rays ℓ', ℓ'' each of which is allowed to be either horizontal or vertical. Moreover, in the case of a horizontal ray, we assume that its y-coordinate lies in $\mathbb{Z} \pm \frac{1}{10}$, and in the case of a vertical ray we assume that its x-coordinate is 0.

Below we will mainly work with the following two types of such curves. The first one is $\gamma^\uparrow = \{x = 0\} \subset \mathbb{R}^2$ (i.e., the x-axis with its standard orientation) and we take $f_{\gamma^\uparrow} \equiv 0$. Then for every exact Lagrangian $L \subset X$ we can identify $f_{\gamma^\uparrow,L}$ with f_L in the obvious way.

The second type is the curve $\gamma_{i,j}$, where $i \leq j$ are two integers, depicted in Figure 3.11 and oriented by going from the lower horizontal end to the upper horizontal end. Note that by taking $\gamma_{i,j}$ close enough to the dotted polygonal curve in Figure 3.11 we can assume that $\lambda_{\mathbb{R}^2}|_{\gamma_{i,j}}$ is very close to 0. We fix the primitive $f_{\gamma_{i,j}}$ to be the one that vanishes on the vertical part of $\gamma_{i,j}$.

Let $\gamma \subset \mathbb{R}^2$ and f_γ be as above. Following [7, Section 4.2] and [10, Section 3.6], there is a (weakly) filtered A_∞-functor $\mathcal{I}_\gamma : W\mathcal{F}uk(X) \to W\mathcal{F}uk(\mathbb{R}^2 \times X)$, called an inclusion functor, which sends the object $L \in \mathrm{Obj}(W\mathcal{F}uk(X))$ to the object $\gamma \times L \in \mathrm{Obj}(\mathcal{F}uk(\mathbb{R}^2 \times X))$. The first-order component of $\mathcal{I}_\gamma$ is a chain map

$$(\mathcal{I}_\gamma)_1 : \mathrm{CF}(N_0, N_1) \to \mathrm{CF}(\gamma \times N_0, \gamma \times N_1),$$

defined for all exact Lagrangians N_0, N_1, which induces an isomorphism in homology. Note that, since $\gamma \times N_0$ and $\gamma \times N_1$ do not intersect transversely (unless we have $N_0 \cap N_1 = \emptyset$), we need to use here Floer data with non-trivial Hamiltonians that also involve a component in the $\mathbb{R}^2$-direction. We skip these details here and refer the reader to [7, Section 4.2] and [10, pp. 68–69] for the precise details. The higher-order components $(\mathcal{I}_\gamma)_d$, $d \geq 2$, of $\mathcal{I}_\gamma$ are defined to be 0.

[2]The orientation on γ is necessary in order to set up a graded Floer theory and also to work with coefficients over rings of characteristic $\neq 2$. Here we work with $\mathbb{Z}_2$-coefficients; therefore, if one wants to ignore the grading then the orientation of γ becomes irrelevant.

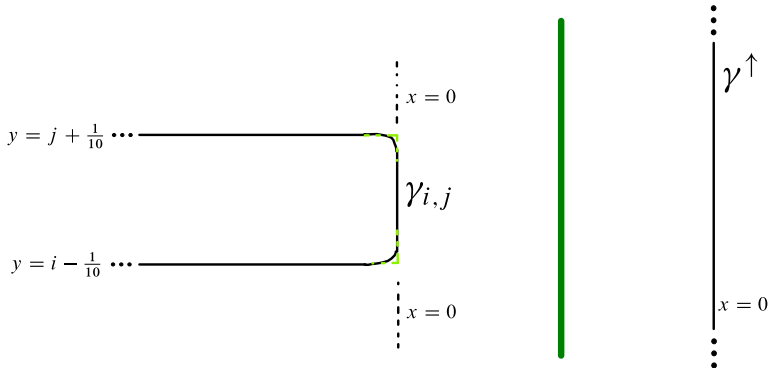


Figure 3.11. The curves $\gamma_{i,j}$ and $\gamma^{\uparrow}$.

Let $V \subset \mathbb{R}^2 \times X$ be a Lagrangian cobordism. Denote by $\mathcal{V}$ the (weakly) filtered Yoneda module of V and consider the pullback module $\mathcal{I}_{\gamma}^{*}\mathcal{V}$. Note that for every exact Lagrangian $N \subset X$ we have

$$\mathcal{I}_{\gamma}^{*}\mathcal{V}(N) = \mathrm{CF}(\gamma \times N, V)$$

as *filtered chain complexes*. (The filtrations are induced by f_V, $f_{\gamma,N}$, and by the Floer data in case it is not trivial.)

Assume that the ends of V are $L_1, \ldots, L_k$ and, moreover, that V is cylindrical over $(-\infty, \delta] \times \mathbb{R}$ for some $\delta > 0$. (This can always be achieved by a suitable translation along the x-axis.) Fix $1 \leq i \leq k-1$ and consider the curve $\gamma_{i,i+1}$ and the pullback (weakly) filtered module $\mathcal{I}_{\gamma_{i,i+1}}^{*}\mathcal{V}$. The cobordism V gives rise to a module homomorphism

$$\Gamma_{V,\gamma_{i,i+1}} : L_{i+1} \to L_i$$

which preserves action filtrations and such that

$$\mathcal{I}_{\gamma_{i,i+1}}^{*}\mathcal{V} = T^{d_i}\mathrm{Cone}(L_{i+1} \xrightarrow{\Gamma_{V,\gamma_{i,i+1}}} L_i) \tag{3.36}$$

as (weakly) filtered modules. Here T stands for the grading-translation functor, and the amount $d_i \in \mathbb{Z}$ by which we translate depends only on i. We will be more precise about the values of d_i later on. The references for the construction of the map $\Gamma_{V,\gamma_{i,i+1}}$ are the following. For the unfiltered case, see [7, Section 4.4 and Proposition 4.4.1]. The map $\Gamma_{V,\gamma_{i,i+1}}$ is constructed on page 1805 of that paper, where it is denoted by ϕ_j; see also Proposition 4.4.3 there. The weakly filtered case is treated in [10]; see Proposition 3.5 and its proof, pp. 73–76. More relevant background material on inclusion functors and iterated cones can be found in Sections 3.6 and 3.7 of that paper. Note that here we are working in a strictly filtered setting (which is a

special case of the weakly filtered case), and this simplifies many of the arguments from [10]. In addition to these references, we provide below in Section 3.5.1.4 an outline of the construction of $\Gamma_{V,\gamma_{i,i+1}}$, avoiding technicalities.

The map $\Gamma_{V,\gamma_{i,i+1}}$ is canonically defined by V and $\gamma_{i,i+1}$, up to a boundary in the chain complex $\hom^{\leq 0}_{\mathrm{mod}_{\mathcal{F}uk(M)}}(L_{i+1}, L_i)$. Therefore it gives rise to a well-defined morphism in the homological persistence category $\mathcal{C}_0 = H\big(\hom^{\leq 0}_{\mathrm{mod}_{\mathcal{F}uk(M)}}(L_{i+1}, L_i)\big)$, which by abuse of notation we still denote by $\Gamma_{V,\gamma_{i,i+1}} \in \hom_{\mathcal{C}_0}(L_{i+1}, L_i)$.

The above can be generalized to several consecutive ends in a row as follows. Fix $1 \leq i \leq j \leq k$. The pullback module $\mathcal{I}^*_{\gamma_{i,j}}\mathcal{V}$ can be identified with an iterated cone of the type

$$\mathcal{I}^*_{\gamma_{i,j}}\mathcal{V} = \mathrm{Cone}\big(L_j \longrightarrow \mathrm{Cone}(L_{j-1} \longrightarrow \ldots \longrightarrow \mathrm{Cone}(L_{i+1} \longrightarrow L_i)\ldots)\big), \quad (3.37)$$

where, similarly to the case $\Gamma_{V,\gamma_{i,i+1}}$, all the maps in the iterated cone are module homomorphisms that preserve filtrations. See Figure 3.12. The references given above for the construction of $\Gamma_{V,\gamma_{i,i+1}}$ are relevant also for the construction of (3.37).

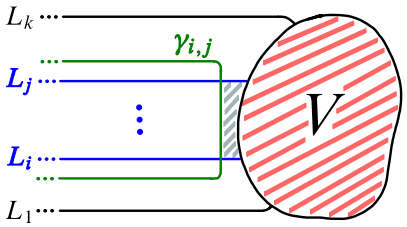


Figure 3.12. The module $\mathcal{I}^*_{\gamma_{i,j}}\mathcal{V}$.

Note that there are some grading translations in (3.37) that we have ignored. We will be more precise about this point later on, when we consider iterated cones involving three objects.

Remark 3.33. If V has k ends, then for every $i \leq 1$ and $k \leq l$ we have $\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V} = \mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V}$.

Finally, fix $1 \leq i \leq l < j \leq k$ and consider the two modules $\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}$ and $\mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V}$. There is a module homomorphism $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}} : \mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V} \to \mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}$ which preserves filtrations. Note that we have

$$\mathcal{I}^*_{\gamma_{i,j}}\mathcal{V} = T^{d_{i,l,j}}\,\mathrm{Cone}\big(\mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V} \xrightarrow{\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}} \mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}\big)$$

for some $d_{i,l,j} \in \mathbb{Z}$. See Figure 3.13. While the construction of $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}$ does not explicitly appear in the references mentioned after (3.36), it can be easily deduced from the material in those papers. See also Section 3.5.1.4 below.

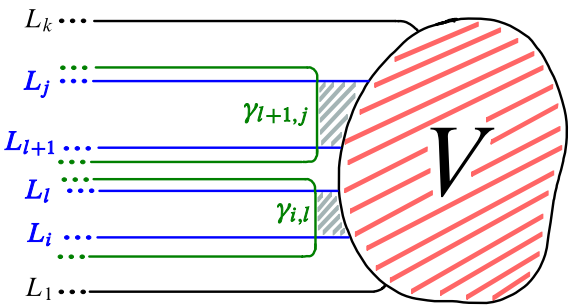


Figure 3.13. The modules $\mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V}$ and $\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}$.

3.5.1.4 Module maps induced by cobordisms. The purpose of this section is to outline the constructions of the module maps $\Gamma_{V,\gamma_{i,i+1}} : L_{i+1} \to L_i$ and $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}} : \mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V} \to \mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}$ from Section 3.5.1.3. We will not give a fully rigorous account of the subject here, in an attempt to avoid technicalities as much as possible. Full details can be found in the references given after (3.36).

We begin with the map $\Gamma_{V,\gamma_{i,i+1}} : L_{i+1} \to L_i$. We will first explain how to construct the first-order component $(\Gamma_{V,\gamma_{i,i+1}})_1$ of this map.

Consider the pullback module $\mathcal{M} := \mathcal{I}^*_{\gamma_{i,i+1}}\mathcal{V}$ and Figure 3.14. For every exact Lagrangian N we have the following equalities of vector spaces:

$$\mathcal{M}(N) = \mathrm{CF}(\gamma_{i,i+1} \times N, V) = \mathrm{CF}(N, L_{i+1}) \oplus \mathrm{CF}(N, L_i). \tag{3.38}$$

In terms of Figure 3.14, the first summand corresponds to the intersection points $N \cap L_{i+1}$ lying above the point P, and the second summand to the intersection points $N \cap L_i$ lying above Q. For the sake of illustration, we have made here several simplifying assumptions (which cannot really be made in general). Namely, that N intersects both L_i and L_{i+1} transversely and that we can take the Floer and perturbation data for $W\mathcal{F}uk(\mathbb{R}^2 \times X)$ and $W\mathcal{F}uk(X)$ to have 0 Hamiltonian terms.

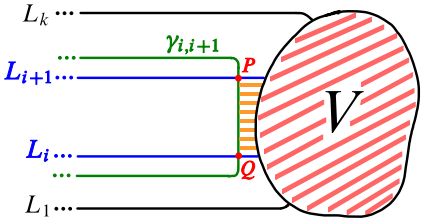


Figure 3.14. The module map $\Gamma_{V,\gamma_{i,i+1}} : L_{i+1} \to L_i$. The projection to $\mathbb{R}^2$ of the strips of type (PQ) is depicted as the orange horizontally-striped region.

Next we consider the differential $\mu_1^{\mathcal{M}}$ of this module. Again, for simplicity, assume that the almost complex structure J in the Floer data for $\mathrm{CF}(\gamma_{i,i+1} \times N, V)$ is chosen so that the projection $\pi : \mathbb{R}^2 \times X \to \mathbb{R}^2$ is (J, i)-holomorphic, where i is the standard complex structure on $\mathbb{R}^2 \cong \mathbb{C}$. To describe $\mu_1^{\mathcal{M}}$ we need to consider Floer strips contributing to the differential of $\mathrm{CF}(\gamma_{i,i+1} \times N, V)$. A priori, these are of four types:

(PP) Strips going from intersection points above P to points above P.

(QQ) Strips going from intersection points above Q to points above Q.

(PQ) Strips going from intersection points above P to points above Q.

(QP) Strips going from intersection points above Q to points above P.

Our assumptions on J and the Hamiltonians in the Floer data imply that all the PP and QQ strips have constant projection to $\mathbb{R}^2$, hence completely lie in $\{P\} \times X$ and $\{Q\} \times X$, respectively. Moreover, when viewed as strips in X, they are in one-to-one correspondence with the Floer strips that contribute to the differentials on $\mathrm{CF}(N, L_{i+1})$ and $\mathrm{CF}(N, L_i)$, respectively.

Standard arguments based on complex analysis in the plane (e.g., the open mapping theorem as used in [6, Section 4], [7, Sections 3-4]) imply that there are no Floer strips of type (QP). However, strips of type (PQ) may definitely exist and we write their contribution to the Floer complex as a linear map: $\phi_1 : \mathrm{CF}(N, L_{i+1}) \to \mathrm{CF}(N, L_i)$, which is based on counting strips emanating from an intersection point above P to an intersection point above Q.

Summing up, the differential $\mu_1^{\mathcal{M}}$ can be written as

$$\mu_1^{\mathcal{M}}(x^P, x^Q) = \big(\mu_1(x^P), \mu_1(x^Q) + \phi_1(x^P)\big).$$

Here, x^P, x^Q are intersection points corresponding to the first and second summands in (3.38), and μ_1 stands for the Floer differentials coming from $W\mathcal{F}uk(X)$. The fact that $\mu_1^{\mathcal{M}}$ is a differential implies that ϕ_1 is a chain map and, moreover, that $\mathcal{M}(N) = \mathrm{Cone}(L_{i+1} \xrightarrow{\phi_1} L_i)$ as chain complexes. We define the first-order component of $\Gamma_{V,\gamma_{i,i+1}}$ to be $(\Gamma_{V,\gamma_{i,i+1}})_1 := \phi_1$.

The construction of the higher-order components of $\Gamma_{V,\gamma_{i,i+1}}$ is similar, though technically more involved. To construct the maps $(\Gamma_{V,\gamma_{i,i+1}})_d$, $d \geq 2$, we need to analyze the μ_d-operations of the module $\mathcal{M}$. Fix d exact Lagrangians $N_0, \ldots, N_{d-1}$ in X and $\underline{y} = (y_1, \ldots, y_{d-1})$ with $y_k \in \mathrm{CF}(N_{k-1}, N_k)$. Since $\mathcal{I}_{\gamma_{i,i+1}}$ has vanishing higher-order components, we have

$$\mu_d^{\mathcal{M}}(\underline{y}, (x^P, x^Q)) = \mu_d^{\mathbb{R}^2 \times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(y_1), \ldots, (\mathcal{I}_{\gamma_{i,i+1}})_1(y_{d-1}), (x^P, x^Q)\big), \tag{3.39}$$

where $\mu_d^{\mathbb{R}^2 \times X}$ is the dth-order A_∞-operation in $W\mathcal{F}uk(\mathbb{R}^2 \times X)$ and the points $x^P \in \mathrm{CF}(N_0, L_{i+1})$, $x^Q \in \mathrm{CF}(N_0, L_i)$ are as before.

The right-hand side of equation (3.39) counts Floer d-polygons with "edges" on the Lagrangians $\gamma_{i,i+1} \times N_0, \ldots, \gamma_{i,i+1} \times N_{d-1}, V$. By arguments similar to

those used for $d = 1$, one shows that there are no Floer polygons with entry points in $(\mathcal{I}_{\gamma_{i,i+1}})_1(y_1), \ldots, (\mathcal{I}_{\gamma_{i,i+1}})_1(y_{d-1}), x_Q$ and exit points lying above P. Consequently, (3.39) has the shape

$$\begin{aligned}\mu_d^{\mathcal{M}}(\underline{y}, (x^P, x^Q)) \\ = \Big(\mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y}), x^P\big), \mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y}), x^Q\big) + \phi_d\big(\underline{y}, x^P\big)\Big)\end{aligned} \tag{3.40}$$

with respect to the splitting $\mathcal{M}(N_0) = \mathrm{CF}(N_0, L_{i+1}) \oplus \mathrm{CF}(N_0, L_i)$ used before. Here $(\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y})$ stands for $\big((\mathcal{I}_{\gamma_{i,i+1}})_1(y_1), \ldots, (\mathcal{I}_{\gamma_{i,i+1}})_1(y_{d-1})\big)$. Thus $\phi_d(\underline{y}, x^P)$ counts Floer polygons in $\mathbb{R}^2 \times X$ with entry points $(\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y})$, x^P and an exit point over Q. The dth-order component $(\Gamma_{V,\gamma_{i,i+1}})_d$ of the desired map $\Gamma_{V,\gamma_{i,i+1}}$ is the multilinear map ϕ_d.

It remains to explain why $\mathcal{M} = \mathcal{I}^*_{\gamma_{i,i+1}}\mathcal{V}$ is the mapping cone (in the A_∞-sense) of the map $\Gamma_{V,\gamma_{i,i+1}}$. Consider the other two terms on the right-hand side of (3.40), namely

$$\mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y}), x^P\big) \quad \text{and} \quad \mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y}), x^Q\big).$$

These two terms can be identified with

$$\mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i+1,i+1}})_1(\underline{y}), x^P\big) \quad \text{and} \quad \mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i}})_1(\underline{y}), x^Q\big),$$

respectively (note we are using now the curves $\gamma_{i+1,i+1}$ and $\gamma_{i,i}$, and not $\gamma_{i,i+1}$). In other words, the preceding two expressions can be identified with the μ_d-operations of the pullback modules $\mathcal{I}^*_{\gamma_{i+1,i+1}}\mathcal{V}$ and $\mathcal{I}^*_{\gamma_{i,i}}\mathcal{V}$ applied to $(\underline{y}, x^P)$ and $(\underline{y}, x^Q)$, respectively.

Now, it follows from [7, Section 4.2] that the pullback modules $\mathcal{I}^*_{\gamma_{i+1,i+1}}\mathcal{V}$ and $\mathcal{I}^*_{\gamma_{i,i}}\mathcal{V}$ are quasi-isomorphic to the Yoneda modules of L_{i+1} and L_i, respectively. Therefore, up to grading translation, we obtain $\mathcal{M} = \mathrm{Cone}(L_{i+1} \xrightarrow{\phi} L_i)$, where $\phi = \{\phi_d\}_{d\geq 1}$, which proves (3.36). This concludes our rough outline of the construction of the map $\Gamma_{V,\gamma_{i,i+1}}$.

The definition of the maps $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}$ is similar to the above and we will just go over the main points of the construction. Put $\mathcal{Q} = \mathcal{I}^*_{\gamma_{i,j}}\mathcal{V}$ and consider Figure 3.15. It is not hard to show that for every exact Lagrangian N we have

$$\mathcal{Q}(N) = \mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V}(N) \oplus \mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}(N) \tag{3.41}$$

as vector spaces. Elements of the first summand can be written as $\underline{x}^P = (x^{P_j}, \ldots, x^{P_{l+1}})$, with $x^{P_j} \in \mathrm{CF}(N, L_k)$ viewed as lying above the point P_k in Figure 3.15. Similarly, elements of the second summand of (3.41) can be written as $\underline{x}^Q = (x^{Q_i}, \ldots, x^{Q_l})$. The differential $\mu_1^{\mathcal{Q}}$ of this module turns out to have the following shape:

$$\mu_1^{\mathcal{Q}}(\underline{x}^P, \underline{x}^Q) = \Big(\mu_1^{\mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V}}(\underline{x}^P), \mu_1^{\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}}(\underline{x}^Q) + \psi_1(\underline{x}^P)\Big),$$

where $\psi_1 : \mathcal{I}^*_{\gamma_{l+1,j}} \mathcal{V}(N) \to \mathcal{I}^*_{\gamma_{i,l}} \mathcal{V}(N)$ is a linear map. The reason for this is similar to what has been explained earlier for the module $\mathcal{M} = \mathcal{I}^*_{\gamma_{i,i+1}} \mathcal{V}$. Namely, there cannot be any Floer strips connecting Q-type points with P-type points. The term $\psi_1(\underline{x}^P)$ counts Floer strips (in $\mathbb{R}^2 \times X$), with one boundary on $\gamma_{i,j} \times N$ and the other boundary on V, emanating from any entry x^{P_k} of $\underline{x}^P$ and going to some entry x^{Q_m} of $\underline{x}^Q$. In terms of Figure 3.15, the projections of, e.g., the strips that go from $x^{P_{l+1}}$ to x^{Q_l} are depicted in light blue (oblique stripes). The projections of the strips corresponding to $\mu_1^{\mathcal{I}_{\gamma_{l+1,j}}{}^* \mathcal{V}}$ are in orange (horizontal stripes), and those corresponding to $\mu_1^{\mathcal{I}^*_{\gamma_{i,l}} \mathcal{V}}$ $\mu_1^{\mathcal{I}^*_{\gamma_{i,l}} \mathcal{V}}$ are in purple (vertical stripes).

The first-order component of the desired map $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}$ is defined to be the map ψ_1. The construction of the higher-order components of $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}$ is analogous to the case of $\Gamma_{V,\gamma_{i,i+1}}$ discussed earlier.

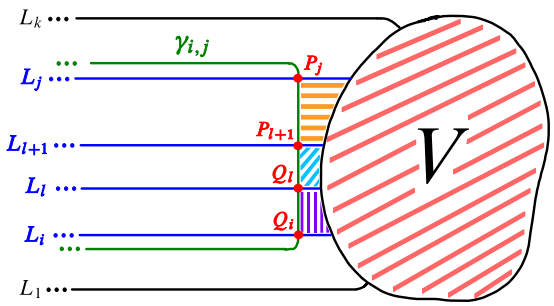


Figure 3.15. The module map $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}} : \mathcal{I}^*_{\gamma_{l+1,j}} \mathcal{V} \to \mathcal{I}^*_{\gamma_{i,l}} \mathcal{V}$.

3.5.1.5 Shadow of cobordisms. Given a Lagrangian cobordism $V \subset \mathbb{R}^2 \times X$, we define its *outline* [24] as

$$\mathrm{out}(V) := \mathbb{R}^2 \setminus \mathcal{U},$$

where $\mathcal{U} \subset \mathbb{R}^2 \setminus \pi(V)$ is the union of all the *unbounded* connected components of $\mathbb{R}^2 \setminus \pi(V)$. An important measurement associated to V is its *shadow* [10, 24], $\mathcal{S}(V)$:

$$\mathcal{S}(V) := \mathrm{Area}(\mathrm{out}(V)).$$

Note that $\mathrm{out}(V) \subset \mathbb{R}^2$ is a measurable set, hence $\mathcal{S}(V)$ is well defined. The shadow plays a central role in defining cobordism-related metrics on spaces of Lagrangians [10].

For the purposes of this section, it will be easier to work with a slightly different variant of the shadow, which we call the *exterior shadow*. Fix a rectangle $Q \subset \mathbb{R}^2$ large enough that $\mathrm{out}(V) \setminus Q$ consists of only horizontal rays, and write

$$\mathrm{out}_Q(V) = Q \cap \mathrm{out}(V).$$

Define the *exterior shadow* of V to be

$$\mathcal{S}_{\mathrm{e}}(V) = \inf\big\{A \mid \exists \text{ a smooth embedding } \varphi : B \to \mathbb{R}^2, \text{ with } \\ \text{image}(\varphi) \supset \mathrm{out}_Q(V) \text{ and } \mathrm{Area}(\mathrm{image}(\varphi)) \geq A\big\}.$$

Here $B \subset \mathbb{R}^2$ stands for the closed 2-dimensional unit disk. It is not hard to see that $\mathcal{S}_{\mathrm{e}}(V)$ is independent of the choice of the rectangle Q.

Obviously, we have $\mathcal{S}_{\mathrm{e}}(V) \geq \mathcal{S}(V)$ (because $\mathcal{S}(V) = \mathrm{Area}(\mathrm{out}_Q(V))$). But in fact, we actually have $\mathcal{S}_{\mathrm{e}}(V) = \mathcal{S}(V)$. Since this statement has not been proved in full in the literature (though see [24, p. 33] for a related partial argument), we include in Section 3.5.1.11 below a sketch of a proof that $\mathcal{S}_{\mathrm{e}}(V) = \mathcal{S}(V)$.

Previous papers on the subject used the shadow rather than the exterior shadow. However, for the rest of this section, whose purpose is mainly illustrative, we opt for the exterior shadow since it is more intuitive to work with.

3.5.1.6 r-acyclic objects. Let $V \subset \mathbb{R}^2 \times X$ be a cobordism with ends $L_1, \ldots, L_k$ such that V is cylindrical over $(-\infty, \delta] \times \mathbb{R}$ for some $\delta > 0$. Let $\mathcal{S}_{\mathrm{e}}(V)$ be the exterior shadow of V and denote by σ the area of the region to the right of $\gamma^{\uparrow}$ enclosed between $\gamma^{\uparrow}$ and the projection to $\mathbb{R}^2$ of the non-cylindrical part of V. See Figure 3.16.

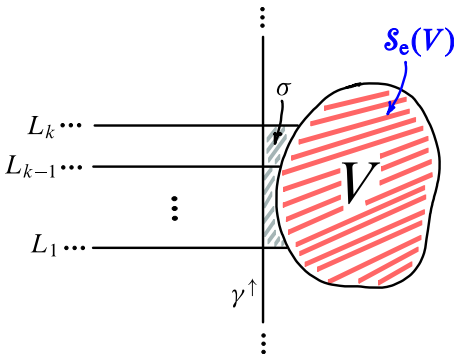


Figure 3.16. The projection of a cobordism to $\mathbb{R}^2$, its (exterior) shadow $\mathcal{S}_{\mathrm{e}}(V)$ and the area σ.

The module $\mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V} = \mathcal{I}^*_{\gamma_{1,k}}\mathcal{V}$ (which also has the description (3.37) with $i = 1$, $j = k$) is r-acyclic, where $r := \mathcal{S}_{\mathrm{e}}(V) + \sigma$. This can be easily seen from the fact that V can be deformed to a cobordism W which is disjoint from $\gamma^{\uparrow} \times X$, via a compactly supported Hamiltonian isotopy whose Hofer length is at most r. Standard Floer theory then implies that $\mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V}$ is acyclic of boundary depth at most r. In the terminology used in this memoir, this means that the object $\mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V}$ is r-acyclic.

Remark 3.34. The area summand σ that adds to the exterior shadow of V in the quantity r can be made arbitrarily small at the expense of applying appropriate shifts

to each of the ends L_i of V. One way to do this is to replace the curve $\gamma^\uparrow$ by a curve $\gamma_V^\uparrow$ that coincides with $\gamma^\uparrow$ outside a compact subset and such that $\gamma_V^\uparrow$ approximates the shape of the projection of the non-cylindrical part of V in such a way that the area σ' enclosed between $\gamma_V^\uparrow$ and $\pi(V)$ is small. See Figure 3.17. One can apply a similar modification to the curves $\gamma_{i,j}$. Note that, in contrast to $f_{\gamma^\uparrow}$, the primitive $f_{\gamma_V^\uparrow}$ can no longer be assumed to be 0 (a similar remark applies to the primitives of the modifications of $\gamma_{i,j}$). As a result, the cone decompositions (3.37) associated to the pullback modules $\mathcal{I}^*_{\gamma_{i,j}}$ will have the same shape but each of the Lagrangians $L_i, \dots, L_j$ will gain a different shift in action. Note that this will also result in "tighter" weighted exact triangles than the ones we obtain below, in terms of the weights and the various shifts on the objects forming these triangles.

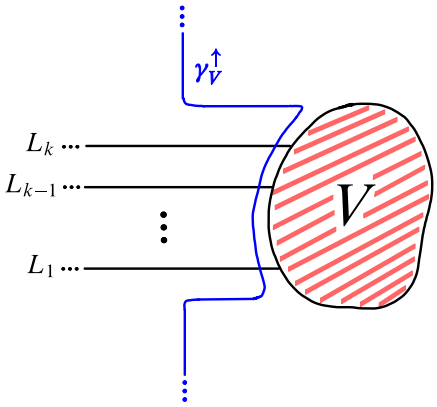


Figure 3.17. Replacing $\gamma^\uparrow$ by a curve $\gamma_V^\uparrow$ better approximating the shape of V.

To simplify the exposition, below we will not make these modifications and will stick to the curves $\gamma^\uparrow$ and $\gamma_{i,j}$ as defined above, at the expense of σ not being necessarily small and the weights of the triangles not being necessarily optimal.

3.5.1.7 *r*-isomorphisms. We begin by visualizing the canonical map $\eta_r^L : \Sigma^r L \to L$, where $L \subset X$ is an exact Lagrangian. Consider the curve $\gamma \subset \mathbb{R}^2$ depicted in Figure 3.18, and let r be the area enclosed between γ and $\gamma^\uparrow$. Let f_γ be the unique primitive of $\lambda_{\mathbb{R}^2}|_\gamma$ that vanishes along the lower end of γ. Note that $f_\gamma \equiv r$ along the upper end of γ. Let $V = L \times \gamma$ and set $f_V := f_{\gamma,L}$. Therefore, the primitives induced by V on its ends are as follows: the primitive on the lower end coincides with f_L, while the primitive on the upper end coincides with $f_L + r$. In other words, the cobordism V has ends L and $\Sigma^r L$. Moreover, the map $\Gamma_{V,\gamma_{1,2}} : \Sigma^r L \to L$ induced by V and $\gamma_{1,2}$ is precisely η_r^L.

Another source of geometric r-isomorphisms comes from Hamiltonian isotopies. Let ϕ_t^H, $t \in [0,1]$, be a Hamiltonian isotopy and let $L \subset X$ be an exact Lagrangian.

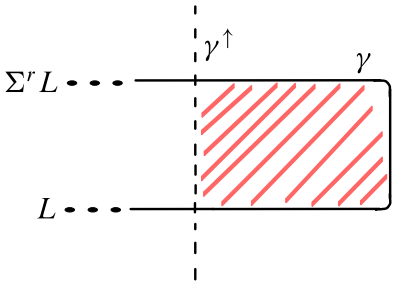


Figure 3.18. The cobordism inducing $\eta_r^L : \Sigma^r L \to L$.

The Lagrangian suspension construction gives rise to an exact Lagrangian cobordism between L and $\phi_1^H(L)$. After bending the ends of that cobordism so that they become negative, one obtains a Lagrangian cobordism whose negative ends are L and $\phi_1^H(L)$. See Figure 3.19. The primitive f_V on V is uniquely defined by the requirement that f_V coincides with f_L on the lower end of V. The exterior shadow $\mathcal{S}_{\mathrm{e}}(V)$ of this cobordism equals the Hofer length of the isotopy $\{\phi_t^H\}_{t\in[0,1]}$, and we obtain an r-isomorphism $\phi_1^H(L) \to L$.

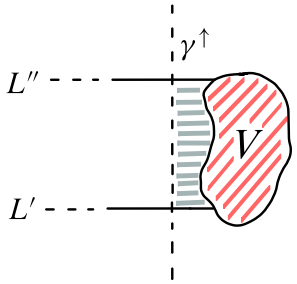


Figure 3.19. The Lagrangian suspension of a Hamiltonian isotopy, after bending the ends.

More generally, let V be a Lagrangian cobordism with ends $L_1, \dots, L_k$. Let $r = \mathcal{S}_{\mathrm{e}}(V) + \sigma$ as above. Fix $1 \le l < k$. As explained above we have, up to an overall translation in grading,

$$\mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V} = \mathcal{I}^*_{\gamma_{1,k}}\mathcal{V} = \mathrm{Cone}\big(\mathcal{I}^*_{\gamma_{l+1,k}}\mathcal{V} \xrightarrow{\Gamma_{V,\gamma_{1,l},\gamma_{l+1,k}}} \mathcal{I}^*_{\gamma_{1,l}}\mathcal{V}\big),$$

and since $\mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V}$ is r-acyclic, the map $\Gamma_{V,\gamma_{1,l},\gamma_{l+1,k}} : \mathcal{I}^*_{\gamma_{l+1,k}}\mathcal{V} \to \mathcal{I}^*_{\gamma_{1,l}}\mathcal{V}$ is an r-isomorphism.

3.5.1.8 Weighted exact triangles. Let $V \subset \mathbb{R}^2 \times X$ be a Lagrangian cobordism with ends $L_1, \dots, L_k$. Let $V' \subset \mathbb{R}^2 \times X$ be the cobordism obtained from V by bending the

upper end L_k clockwise around V so that it goes beyond the end L_1, as in Figure 3.20. To obtain a cobordism according to our conventions, we need to further shift V' upwards by one so that its lower end has y-coordinate 1 (instead of 0). Clearly, V' is also exact and $\mathcal{S}_{\mathrm{e}}(V') = \mathcal{S}_{\mathrm{e}}(V)$. We fix the primitive $f_{V'}$ for V' to be the unique one that coincides with f_V on the ends $L_1, \ldots, L_{k-1}$. A simple calculation shows that $f_{V'}$ induces on the lowest end of V' the primitive $f_{L_k} - r$, where $r = \mathcal{S}_{\mathrm{e}}(V) + \sigma + \epsilon$. (Here ϵ can be assumed to be arbitrarily small. It can be estimated from above by the area enclosed by the bent end corresponding to ℓ_k, the projection to $\mathbb{R}^2$ of the non-cylindrical part of V, ℓ_1, and $\gamma^{\uparrow}$. See Figure 3.20.)

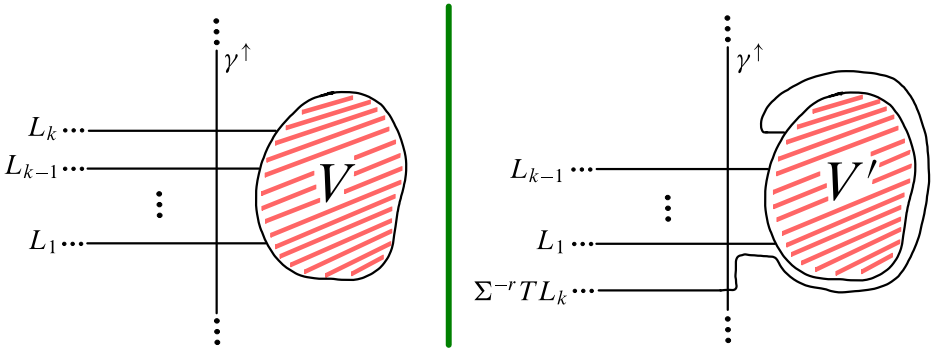


Figure 3.20. Bending the upper end L_k of a cobordism V.

Taking grading into account (in case V is graded in the sense of Floer theory), one can easily see that the grading on the lowest end of V' is translated by 1 in comparison to L_k. Summing up, the above procedure transforms a cobordism V with ends $L_1, \ldots, L_k$ into a cobordism V' with ends $\Sigma^{-r}TL_k, L_1, \ldots, L_{k-1}$.

Similarly, one can take V and bend its lowest end L_1 counterclockwise around V and obtain a new cobordism V'' with ends $L_2, \ldots, L_k, \Sigma^r T^{-1} L_1$ and with $\mathcal{S}_{\mathrm{e}}(V'') = \mathcal{S}_{\mathrm{e}}(V)$.

We are now in a position to describe geometrically weighted exact triangles. Let V be a cobordism with three ends, which for compatibility with Definition 2.42 we denote by C, B, A (going from the lowest end upward). See Figure 3.21.

Let $r = \mathcal{S}_{\mathrm{e}}(V) + \sigma + \epsilon$. Put

$$\bar{u} := \Gamma_{V,\gamma_{2,3}} : A \to B, \quad \bar{v} := \Gamma_{V,\gamma_{1,2}} : B \to C.$$

Consider also the counterclockwise rotation V'' of V whose ends are $B, A, \Sigma^r T^{-1} C$. Let

$$\bar{w} := \Sigma^{-r} T \Gamma_{V'',\gamma_{2,3}} = \Gamma_{V',\gamma_{1,2}} : C \to \Sigma^{-r} TA.$$

We claim that

$$A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r} TA \tag{3.42}$$

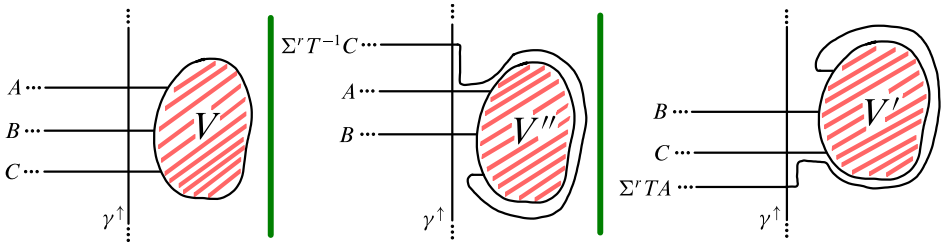


Figure 3.21. A cobordism leading to an exact triangle of weight r.

is a strict exact triangle of weight r. This triangle is based on the genuinely exact triangle from $\mathcal{C}_0$:

$$A \xrightarrow{u} B \xrightarrow{v} C' \xrightarrow{w} TA,$$

where $u = \bar{u}$, $C' = \mathcal{I}^*_{\gamma_{2,3}} V = \mathrm{Cone}(A \xrightarrow{u} B)$, $v : B \to C'$ is the standard inclusion, and $w : C' \to TA$ the standard projection. The r-isomorphism $\phi : C' \to C$ and its right r-inverse $\psi : \Sigma^r C \to C'$ are as follows:

$$\phi = \Gamma_{V,\gamma_{1,1},\gamma_{2,3}}, \quad \psi = T\Gamma_{V'',\gamma_{1,2},\gamma_{3,3}} : \Sigma^r C \to T\mathcal{I}^*_{\gamma_{1,2}} V'' = C'.$$

Note that $\mathcal{I}^*_{\gamma_{1,2}} V'' = T^{-1}\mathcal{I}^*_{\gamma_{2,3}} V = T^{-1}C' = T^{-1}\mathrm{Cone}(A \xrightarrow{u} B)$.

The fact that ψ is a right r-inverse of ϕ and that these maps fit into the diagram (2.20) follows from standard arguments in Floer theory. Note that these statements do not hold on the chain level, but only in $\mathcal{C}_0$.

3.5.1.9 Rotation of triangles. Let V be a cobordism with three ends C, B, A as in Section 3.5.1.8, and consider the exact triangle (3.42) of weight r. Let V' be the clockwise rotation of V, with ends $B, C, \Sigma^{-r}TA$. The exact triangle associated to V' is

$$B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r}TA \xrightarrow{\bar{u}'} \Sigma^{-r-\epsilon'}TB, \tag{3.43}$$

where ϵ' can be assumed to be arbitrarily small. It is not hard to see that in $\mathcal{C}_\infty$ (up to identifying objects with their shifts and ignoring signs in the maps) the exact triangle (3.43) is precisely the rotation of the exact triangle corresponding to (3.42) in $\mathcal{C}_\infty$.

The above shows that rotation of weighted exact triangles coming from cobordisms with three ends preserves weights (up to an arbitrarily small error). Interestingly, this is sharper than the case in a general TPC, described in Proposition 2.46, where the weight of a rotated triangle doubles. See also Remark 2.48.

3.5.1.10 Weighted octahedral property. The weighted octahedral formula from Proposition 2.49 admits too a geometric interpretation in the realm of cobordisms. We will not give the details of this construction here. Instead, we will briefly explain

the cobordism counterpart of cone refinement and why it behaves additively with respect to weights, as described algebraically in Proposition 2.55. Note that weighted cone refinement is one of the main corollaries of the weighted octahedral property.

For simplicity, we focus here on the case described in Example 2.56 and ignore the grading translation T. Assume that we have two cobordisms: V with ends X, B, A, and U with ends A, F, E.

These cobordisms induce two weighted exact triangles of weights $r = \mathcal{S}_e(V) + \sigma_V + \epsilon$ and $s = \mathcal{S}_e(U) + \sigma_U + \epsilon$. By gluing the two cobordisms along the ends corresponding to A we obtain a new cobordism W with four ends X, B, F, E. See Figure 3.22. By the previous discussion, this exhibits X as an iterated cone with linearization (B, F, E), which corresponds precisely to the algebraic cone refinement of X with linearization (B, A) by A with linearization (F, E).

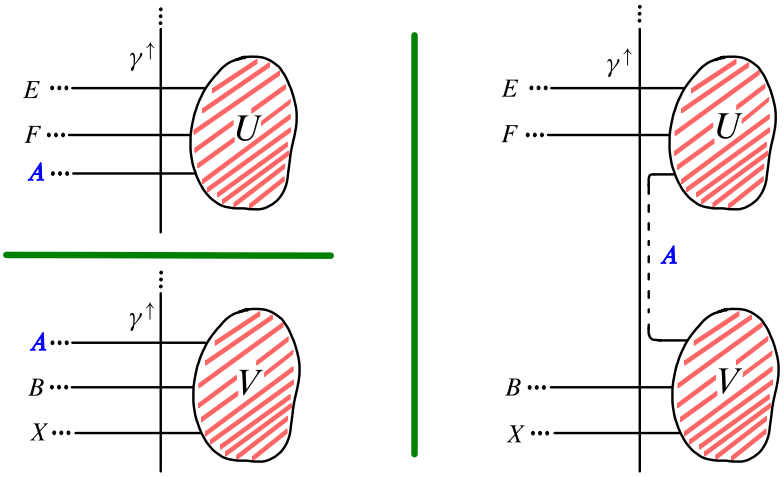


Figure 3.22. Cone refinement via gluing cobordisms along two ends.

Clearly the exterior shadow $\mathcal{S}_e(W)$ of W equals $\mathcal{S}_e(U) + \mathcal{S}_e(V)$, and the total weight of the cone decomposition of X associated to W is $r + s$. Note again that, by the procedure explained in Remark 3.34, one can make the errors e_V, e_U small at the expense of applying some shifts to the elements of the linearization.

Remark 3.35. It is well known that Lagrangian cobordism gives rise to a category with objects Lagrangian submanifolds and with morphisms certain Lagrangian cobordisms (see [7]). Combined with the discussion above, it is natural to wonder whether, by taking into account also the shadows of cobordisms, this category is naturally endowed with a TPC structure. The difficulty in achieving this is that one needs to have a triangulated structure that serves as the level 0 part of this expected TPC. Achieving geometrically such a triangulated structure is delicate, as it requires includ-

ing immersed Lagrangians and cobordisms in the construction. To further produce a TPC structure, this construction needs to be combined with control of cobordism shadows, which makes the whole machinery even more complex. Only partial results in this direction are available at the moment, as in [9].

3.5.1.11 Exterior shadow equals shadow. Here we sketch a proof showing that $\mathcal{S}_{\mathrm{e}}(V) = \mathcal{S}(V)$. Throughout the proof below we will assume that V is connected and that the number of its ends is $l \geq 1$. (In any case, we do not consider cobordisms with no ends in this memoir.)

As at the beginning of 3.5.1.5, fix an open rectangle $Q \subset \mathbb{R}^2$ such that $\mathrm{out}(V) \setminus Q$ consists of only horizontal rays. Denote by $C_1, \ldots, C_l \subset \mathbb{R}^2$ the connected components of $\mathbb{R}^2 \setminus \mathrm{out}(V)$, ordered in counterclockwise order (e.g., with respect to a large enough circle surrounding Q). Set also $C_{Q,i} := C_i \cap Q$, $i = 1, \ldots, l$.

We first claim that each C_i is simply connected. Indeed, this follows easily from the fact that a connected open subset of $\mathbb{R}^2$ is simply connected if and only if every connected component of its complement is unbounded. It follows from the definition of $\mathrm{out}(V)$ that every connected component of $\mathbb{R}^2 \setminus C_i$ is unbounded, hence C_i is simply connected.

Since the C_i's are simply connected, so are the sets $C_{Q,i}$, $i = 1, \ldots, l$. It follows (e.g., by uniformization) that each of the open sets $C_{Q,i}$ is diffeomorphic to an open disk. Furthermore, by the Greene–Shiohama theorem [34] it follows that each $C_{Q,i}$ is in fact symplectomorphic to an open disk $\mathrm{Int}\, B^2(R_i)$ of some radius R_i, endowed with its standard symplectic structure. Fix such symplectomorphisms $\psi : B^2(R_i) \to C_{Q,i}$ for all i.

Assume for the moment that $l \geq 2$ (note that l is precisely the number of ends of V). Reduce the radii R_i slightly to $R_i' = R_i - \delta$ for small $\delta > 0$, and consider the corresponding domains $C'_{Q,i} = \psi(B^2(R_i'))$. Next, connect the boundary of $C'_{Q,i}$ to the boundary of $C'_{Q,i+1}$ by a small strip J_i that intersects $\mathrm{out}_Q(V)$ only along the areas where $\mathrm{out}_Q(V)$ consists solely of horizontal rays. If we smoothly (not symplectically) parametrize J_i as $[-\epsilon, \epsilon] \times [0, 1]$, we just embed J_i in $\mathbb{R}^2$ in such a way that:

(1) The area of J_i is very small and $J_i \subset Q$.

(2) $[-\epsilon, \epsilon] \times \{0\}$ is mapped to $\partial C'_{Q,i}$ near one of the horizontal rays, say E_i, of $\mathrm{out}_Q(V)$ that lies near $\partial C'_{Q,i}$.

(3) $[-\epsilon, \epsilon] \times \{1\}$ is mapped to $\partial C'_{Q_{i+1}}$ near the same horizontal ray E_i we have just used in (1) above.

(4) The rest of J_i intersects $\mathrm{out}(V)$ only along E_i.

We can think of the outcome of connecting $C'_{Q,i}$ to $C'_{Q,i+1}$ with J_i as the boundary connected sum of the closures of the domains $C'_{Q,i}$ and $C'_{Q,i+1}$.

We perform the above construction for all $1 \leq i \leq l-1$, and finally we connect $C'_{Q,l}$ back to $C'_{Q,1}$ in a similar manner, keeping the counterclockwise direction.

Denote by C'_Q the union of all the domains $C'_{Q,i}$ together with the connecting small strips J_i. The outcome $C'_Q \subset Q$ is a domain diffeomorphic to an annulus. Its inner boundary encircles a domain F which is diffeomorphic to the 2-dimensional disk B, and $\mathrm{out}(V) \setminus F$ consists of only horizontal rays. Moreover, by taking the parameter δ small enough we can assume that $\mathrm{Area}(Q \setminus F) = \mathrm{Area}(Q) - \mathrm{Area}(F)$ is arbitrarily close to $\mathrm{Area}(Q) - \mathrm{Area}(\mathrm{out}_Q(V)) = \mathrm{Area}(Q) - \mathcal{S}(V)$. It follows that $\mathrm{Area}(F)$ is arbitrarily close to $\mathcal{S}(V)$ and, at the same time,

$$\mathrm{Area}(F) \geq \mathcal{S}_{\mathrm{e}}(V) \geq \mathcal{S}(V).$$

This concludes the proof under the assumption that $l \geq 2$.

The case $l = 1$ is very similar, only that now we have just one domain $C'_{Q,1}$ and we form the annulus C'_Q by just gluing the small strip J to connect two portions of the boundary of the same domain $C'_{Q,1}$. ∎

3.5.2 Some explicit estimates

We will illustrate here the statements in Theorem 3.4.

Our base manifold will be denoted here by W, and it is the plumbing of two copies of the disk cotangent bundle D^*S^1 of S^1, as in Figure 3.23. The family $\mathcal{F}$ has two elements, F_1 and F_2, as in this figure. They intersect at the single point P.

The primitives on both F_1 and F_2 are the functions identically equal to 0. The family $\mathcal{X}$ consists of F_1, F_2, and the Lagrangians Y, Z, X, and N from Figures 3.24 and 3.25.

The Lagrangian Y is constructed from the surgery $F_2 \# F_1$ at the point P (with a small handle), followed by a small Hamiltonian perturbation. It is easy to see that, for Y to be exact, the "small" gray triangle STP must have the same area as the "large" triangle with the same vertices (only the corners of the second triangle are grayed in the figure). We will denote the area of these triangles by A_Y.

Similarly, the Lagrangian X is constructed from the surgery $F_1 \# F_2$ at the point P (again with a small handle), followed by a small Hamiltonian perturbation. Again, the "small" gray triangle QRP has the same area as the "large" triangle with the same vertices, and we denote this area by A_X. The Lagrangian N is obtained from F_2 by a Hamiltonian perturbation that is large: its Hofer distance equals the area of the strip comprised between N and F_2 and the points x_1 and x_2 (there are two such strips, but they both have the same area). We will denote this area by A_N. Similarly, the Lagrangian Z is obtained from F_1 by a large Hamiltonian perturbation.

The first obvious remark is that X and Y are not smoothly isotopic, because homologically $[X] = [F_2] - [F_1]$ and $[Y] = [F_2] + [F_1]$, and these are not equal in $H_1(X, \mathbb{Z})$.

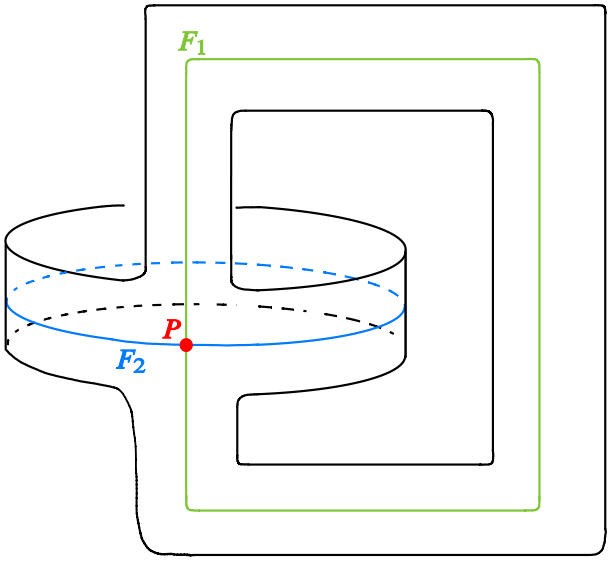


Figure 3.23. The manifold W and the Lagrangians F_1 and F_2.

We are interested in the distance $D^{\mathcal{F}}$.

Lemma 3.36. *We have the inequalities*

$$\frac{A_X}{4} \leq D^{\mathcal{F}}(X,0) \leq 2A_X \frac{A_Y}{4} \leq D^{\mathcal{F}}(Y,0) \leq 2A_Y, \tag{3.44}$$

$$\frac{\max\{A_Y, A_X\}}{4} \leq D^{\mathcal{F}}(X,Y) \leq 2A_X + 2A_Y. \tag{3.45}$$

Proof. We first show the upper bounds in (3.44). The cases of X and Y are entirely similar, and we focus on X. For this, we consider the cone of the map $F_1 \xrightarrow{P} F_2 \to K_1$ (constructed in terms of A_∞-Yoneda modules). We claim that the module K_1 can be mapped to the Yoneda module of X by a quasi-isomorphism. The simplest way to see this geometrically is the following: interpret the module K_1 as the Yoneda module of a marked immersed Lagrangian with one marked self-intersection point (marked in the order (F_1, F_2)). This type of Lagrangians is discussed in [9], for instance. The map we are looking for is of the form $\psi = \mu_2(-, R) : K_1 \to X$, with R the intersection point in Figure 3.25. Of course, once we "guess" this morphism, we can write it purely algebraically. It is easy to see that this is a quasi-isomorphism. For instance, applying it to the Lagrangian N in the picture, it sends x_i to y_i for $i = 1, 2$. Moreover, there is also a quasi-isomorphism $\phi = \mu_2(-, Q) : X \to K_1$, which is a quasi-inverse of the first (on N it is an actual inverse). We can fix the primitive on X that vanishes at the point Q, and thus the primitive on X has value A_X at R. In the terminology of

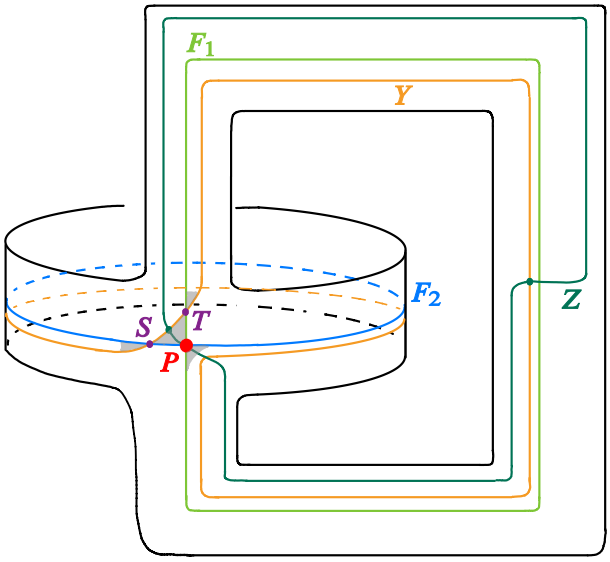


Figure 3.24. The Lagrangian Y obtained as a small perturbation of the surgery $F_2 \#_P F_1$.

this memoir, we have maps of filtered modules

$$\Sigma^{A_X} K_1 \to X \to K_1$$

whose composition agrees with the map η_{A_X} (in other words, the composition is the identity if the filtration is neglected, but once the filtration is taken into account, it shifts the filtration by A_X). We also have a similar identity in the opposite direction.

By applying the same argument as in the second part of Lemma 2.85, we deduce that ϕ and ψ are $2A_X$-isomorphisms, which implies an inequality for the half-distance $\bar{\delta}^{\mathcal{F}}(X, 0) \leq 2A_X$. The other inequality, for the second half-distance, is easy to obtain using the fact that the cone of $\phi : X \to K_1$ is $2A_X$-acyclic, and this implies our upper bound.

For the lower bound, notice that $\delta(X; F_1 \cup F_2) = 2A_X$, and thus the lower bound follows from Theorem 3.4 (ii) (here $\delta(-;-)$ is the relative Gromov width as in Section 3.1).

Clearly, in an entirely similar way we also have $D^{\mathcal{F}}(Y, 0) \leq 2A_Y$ and thus

$$D^{\mathcal{F}}(X, Y) \leq 2A_X + 2A_Y,$$

which is the upper bound in (3.45).

Remark 3.37. The first part of the argument is very similar to the one relating the spectral distance to the distance $D^{\mathcal{F}}$. Indeed, one can think about the two points R and Q as representing the point class and the fundamental class in $\mathrm{HF}(X, K_1)$, and then the first point of Theorem 3.4 implies $D^{\mathcal{F}}(X, K_1) \leq 4A_X$, which means that

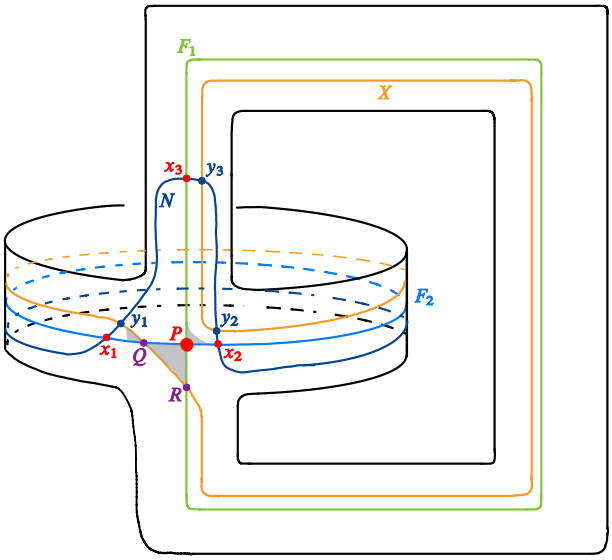


Figure 3.25. The Lagrangian X obtained as a small perturbation of the surgery $F_1 \#_P F_2$ and the Lagrangian N, which is a large Hamiltonian perturbation of F_2.

$D^{\mathcal{F}}(X, 0) \leq 4A_X$ because $D^{\mathcal{F}}(K_1, 0) = 0$. It is very likely that we actually have $D^{\mathcal{F}}(X, 0) = A_X$ and $D^{\mathcal{F}}(Y, 0) = A_Y$.

Finally, we discuss the lower bound in (3.45). For this purpose, we will use here point (ii) of Theorem 3.4. It can be easily shown that $\delta(X; Y \cup F_1 \cup F_2) \geq 2A_X$. Thus we get from point (ii) of Theorem 3.4 that $D^{\mathcal{F}}(X, Y) \geq \frac{A_X}{4}$, as claimed. By symmetry we also get $D^{\mathcal{F}}(X, Y) \geq \frac{A_Y}{4}$. ∎

Remark 3.38. There is an alternative (and possibly more interesting) argument which however gives a slightly weaker inequality than the left-hand side of (3.45). Namely, it implies that

$$\frac{\max\{A_Y, A_X\}}{8} \leq D^{\mathcal{F}}(X, Y).$$

This argument is based on point (iii) of Theorem 3.4 and goes as follows. Consider the Lagrangian Z in Figure 3.24. It has three intersection points with Y and only one with X. By point (iii) of Theorem 3.4, we have $D^{\mathcal{F}}(X, Y) \geq \frac{1}{16}\delta^{\cap}(Z, Y; F_1 \cup F_2)$, where $\delta^{\cap}$ is the quantity defined in (3.2). So this time we need to estimate the number $\delta^{\cap}(Z, Y; F_1 \cup F_2)$. For this estimate it is useful to assume that Z cuts the triangle STP into two pieces of equal area. In this case we have that $\delta^{\cap}(Z, Y; F_1 \cup F_2) = 2A_Y$, and we deduce $D^{\mathcal{F}}(X, Y) \geq \frac{A_Y}{8}$. The inequality involving X follows in the same way, by choosing a deformation Z' of F_1 that this time intersects Y in a single point and X in three points.

References

[1] G. Ambrosioni, Filtered Fukaya categories. *J. Symplectic Geom.* **23** (2025), no. 2, 423–509

[2] T. Asano and Y. Ike, Persistence-like distance on Tamarkin's category and symplectic displacement energy. *J. Symplectic Geom.* **18** (2020), no. 3, 613–649

[3] S. A. Barannikov, The framed Morse complex and its invariants. In *Singularities and bifurcations*, pp. 93–115, Adv. Soviet Math. 21, American Mathematical Society, Providence, RI, 1994

[4] J.-F. Barraud and O. Cornea, Lagrangian intersections and the Serre spectral sequence. *Ann. of Math. (2)* **166** (2007), no. 3, 657–722

[5] U. Bauer and M. Lesnick, Induced matchings and the algebraic stability of persistence barcodes. *J. Comput. Geom.* **6** (2015), no. 2, 162–191

[6] P. Biran and O. Cornea, Lagrangian cobordism. I. *J. Amer. Math. Soc.* **26** (2013), no. 2, 295–340

[7] P. Biran and O. Cornea, Lagrangian cobordism and Fukaya categories. *Geom. Funct. Anal.* **24** (2014), no. 6, 1731–1830

[8] P. Biran and O. Cornea, Bounds on the Lagrangian spectral metric in cotangent bundles. *Comment. Math. Helv.* **96** (2021), no. 4, 631–691

[9] P. Biran and O. Cornea, A Lagrangian pictionary. *Kyoto J. Math.* **61** (2021), no. 2, 399–493

[10] P. Biran, O. Cornea, and E. Shelukhin, Lagrangian shadows and triangulated categories. Astérisque 426, Société Mathématique de France, Paris, 2021

[11] P. Biran, O. Cornea, and J. Zhang, Triangulation and persistence: Algebra 101. 2021, arXiv:2104.12258v1.

[12] A. I. Bondal and M. M. Kapranov, Enhanced triangulated categories (in Russian). *Mat. Sb.* **181** (1990), no. 5, 669–683. English translation: *Math. USSR-Sb.* **70** (1991), no. 1, 93–107

[13] P. Bubenik and N. Milićević, Homological algebra for persistence modules. *Found. Comput. Math.* **21** (2021), no. 5, 1233–1278

[14] G. Carlsson, Topology and data. *Bull. Amer. Math. Soc. (N.S.)* **46** (2009), no. 2, 255–308

[15] F. Charest, *Source spaces and perturbations for cluster complexes*. Ph.D. thesis, Université de Montreal, 2012

[16] F. Charest and C. Woodward, Floer trajectories and stabilizing divisors. *J. Fixed Point Theory Appl.* **19** (2017), no. 2, 1165–1236

[17] F. Chazal, V. de Silva, M. Glisse, and S. Oudot, *The structure and stability of persistence modules*. SpringerBriefs Math., Springer, Cham, 2016

[18] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer, Stability of persistence diagrams. *Discrete Comput. Geom.* **37** (2007), no. 1, 103–120

[19] O. Cornea, Cone-length and Lusternik–Schnirelmann category. *Topology* **33** (1994), no. 1, 95–111

[20] O. Cornea, There is just one rational cone-length. *Trans. Amer. Math. Soc.* **344** (1994), no. 2, 835–848

[21] O. Cornea and F. Lalonde, Cluster homology: An overview of the construction and results. *Electron. Res. Announc. Amer. Math. Soc.* **12** (2006), 1–12

[22] O. Cornea, G. Lupton, J. Oprea, and D. Tanré, *Lusternik–Schnirelmann category*. Math. Surveys Monogr. 103, American Mathematical Society, Providence, RI, 2003

[23] O. Cornea and A. Ranicki, Rigidity and gluing for Morse and Novikov complexes. *J. Eur. Math. Soc. (JEMS)* **5** (2003), no. 4, 343–394

[24] O. Cornea and E. Shelukhin, Lagrangian cobordism and metric invariants. *J. Differential Geom.* **112** (2019), no. 1, 1–45

[25] W. Crawley-Boevey, Decomposition of pointwise finite-dimensional persistence modules. *J. Algebra Appl.* **14** (2015), no. 5, article no. 1550066

[26] G. Dimitrov, F. Haiden, L. Katzarkov, and M. Kontsevich, Dynamical systems and categories. In *The influence of Solomon Lefschetz in geometry and topology*, pp. 133–170, Contemp. Math. 621, American Mathematical Society, Providence, RI, 2014

[27] V. Drinfeld, DG quotients of DG categories. *J. Algebra* **272** (2004), no. 2, 643–691

[28] H. Edelsbrunner, *A short course in computational geometry and topology*. SpringerBriefs Appl. Sci. Technol., Springer, Cham, 2014

[29] H. Edelsbrunner, D. Letscher, and A. Zomorodian, Topological persistence and simplification. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, pp. 454–463, IEEE Computer Society Press, Los Alamitos, CA, 2000

[30] Y.-W. Fan and S. Filip, Asymptotic shifting numbers in triangulated categories. *Adv. Math.* **428** (2023), article no. 109163

[31] Y. Félix, S. Halperin, and J.-C. Thomas, *Rational homotopy theory*. Grad. Texts in Math. 205, Springer, New York, 2001

[32] K. Fukaya and Y.-G. Oh, Zero-loop open strings in the cotangent bundle and Morse homotopy. *Asian J. Math.* **1** (1997), no. 1, 96–180

[33] R. Ghrist, Barcodes: The persistent topology of data. *Bull. Amer. Math. Soc. (N.S.)* **45** (2008), no. 1, 61–75

[34] R. E. Greene and K. Shiohama, Diffeomorphisms and volume-preserving embeddings of noncompact manifolds. *Trans. Amer. Math. Soc.* **255** (1979), 403–414

[35] M. Gromov, *Metric structures for Riemannian and non-Riemannian spaces*. English edn., Mod. Birkhäuser Class., Birkhäuser Boston, Boston, MA, 2007

[36] S. Guillermou, M. Kashiwara, and P. Schapira, Sheaf quantization of Hamiltonian isotopies and applications to nondisplaceability problems. *Duke Math. J.* **161** (2012), no. 2, 201–245

[37] S. Guillermou and P. Schapira, Microlocal theory of sheaves and Tamarkin's non displaceability theorem. In *Homological mirror symmetry and tropical geometry*, pp. 43–85, Lect. Notes Unione Mat. Ital. 15, Springer, Cham, 2014

[38] F. Hensel, Stability conditions and Lagrangian cobordisms. *J. Symplectic Geom.* **18** (2020), no. 2, 463–536

[39] D. Huybrechts, *Fourier–Mukai transforms in algebraic geometry*. Oxford Math. Monogr., The Clarendon Press, Oxford University Press, Oxford, 2006

[40] M. Kashiwara and P. Schapira, *Sheaves on manifolds*. Grundlehren Math. Wiss. 292, Springer, Berlin, 1990

[41] A. Kislev and E. Shelukhin, Bounds on spectral norms and barcodes. *Geom. Topol.* **25** (2021), no. 7, 3257–3350

[42] M. Lesnick, The theory of the interleaving distance on multidimensional persistence modules. *Found. Comput. Math.* **15** (2015), no. 3, 613–650

[43] A. Neeman, *Triangulated categories*. Ann. of Math. Stud. 148, Princeton University Press, Princeton, NJ, 2001

[44] D. Orlov, Triangulated categories of singularities and equivalences between Landau–Ginzburg models (in Russian). *Mat. Sb.* **197** (2006), no. 12, 117–132. English translation: *Sb. Math.* **197** (2006), no. 12, 1827–1840

[45] D. Orlov, Remarks on generators and dimensions of triangulated categories. *Mosc. Math. J.* **9** (2009), no. 1, 143–149

[46] S. Y. Oudot, *Persistence theory: From quiver representations to data analysis*. Math. Surveys Monogr. 209, American Mathematical Society, Providence, RI, 2015

[47] L. Polterovich, D. Rosen, K. Samvelyan, and J. Zhang, *Topological persistence in geometry and analysis*. Univ. Lecture Ser. 74, American Mathematical Society, Providence, RI, 2020

[48] L. Polterovich and E. Shelukhin, Autonomous Hamiltonian flows, Hofer's geometry and persistence modules. *Selecta Math. (N.S.)* **22** (2016), no. 1, 227–296

[49] L. Polterovich, E. Shelukhin, and V. Stojisavljević, Persistence modules with operators in Morse and Floer theory. *Mosc. Math. J.* **17** (2017), no. 4, 757–786

[50] D. Puppe, On the formal structure of stable homotopy theory. In *Colloquium on algebraic topology*, pp. 65–71, Aarhus Universitet Matematisk Institut, 1962.

[51] L. N. Scoccola, *Locally persistent categories and metric properties of interleaving distances*. Ph.D. thesis, The University of Western Ontario, 2020, https://hdl.handle.net/20.500.14721/30236 visited on 15 March 2026

[52] P. Seidel, Graded Lagrangian submanifolds. *Bull. Soc. Math. France* **128** (2000), no. 1, 103–149

[53] P. Seidel, Fukaya categories and Picard–Lefschetz theory. Zur. Lect. Adv. Math., European Mathematical Society (EMS), Zürich, 2008

[54] S. Smale, On the topology of algorithms. I. *J. Complexity* **3** (1987), no. 2, 81–89

[55] D. Sullivan, Infinitesimal computations in topology. *Inst. Hautes Études Sci. Publ. Math.* (1977), no. 47, 269–331

[56] D. Tamarkin, Microlocal condition for non-displaceability. In *Algebraic and analytic microlocal analysis*, pp. 99–223, Springer Proc. Math. Stat. 269, Springer, Cham, 2018

[57] M. Usher, Boundary depth in Floer theory and its applications to Hamiltonian dynamics and coisotropic submanifolds. *Israel J. Math.* **184** (2011), 1–57

[58] M. Usher, Hofer's metrics and boundary depth. *Ann. Sci. Éc. Norm. Supér. (4)* **46** (2013), no. 1, 57–128 (2013)

[59] M. Usher and J. Zhang, Persistent homology and Floer–Novikov theory. *Geom. Topol.* **20** (2016), no. 6, 3333–3430

[60] J.-L. Verdier, Catégories dérivées: Quelques résultats (état 0). In *Cohomologie étale*, pp. 262–311, Lecture Notes in Math. 569, Springer, Berlin, 1977

[61] C. A. Weibel, *An introduction to homological algebra*. Cambridge Stud. Adv. Math. 38, Cambridge University Press, Cambridge, 1994

[62] S. Weinberger, What is... persistent homology? *Notices Amer. Math. Soc.* **58** (2011), no. 1, 36–39

[63] J. Zhang, *Quantitative Tamarkin theory*. CRM Short Courses, Springer, Cham, 2020

[64] A. Zomorodian and G. Carlsson, Computing persistent homology. *Discrete Comput. Geom.* **33** (2005), no. 2, 249–274

MEMOIRS OF THE EUROPEAN MATHEMATICAL SOCIETY

Paul Biran, Octav Cornea, Jun Zhang

Triangulation, Persistence, and Fukaya Categories

This memoir introduces a new algebraic notion: that of a *triangulated persistence category* (TPC), which refines the notion of a triangulated category in the same sense that a persistence module refines that of a vector space. The spaces of morphisms of such a TPC are persistence modules, and the category is endowed with a class of weighted distinguished triangles. Under favourable conditions, we show that the derived Fukaya category admits a TPC refinement, and we apply this to deduce a global rigidity result for spaces of compact, exact Lagrangians in certain Liouville manifolds: we construct a metric on this space with intrinsic symplectic properties.

https://ems.press
ISSN 2747-9080
ISBN 978-3-98547-107-2





