Mathematisches Forschungsinstitut Oberwolfach

Report No. 21/2024

DOI: 10.4171/OWR/2024/21

Applied Harmonic Analysis and Data Science

Organized by
Ingrid Daubechies, Durham
Gitta Kutyniok, München
Holger Rauhut, München

21 April – 26 April 2024

Abstract. Data science is a field of major importance for science and technology nowadays and poses a large variety of challenging mathematical questions. The area of applied harmonic analysis has a significant impact on such problems by providing methodologies both for theoretical questions and for a wide range of applications in machine learning, as well as in in signal and image processing. Building on the success of four previous workshops on applied harmonic analysis in 2012, 2015, 2018, 2021, this workshop focused on several exciting directions, such as mathematical theory of deep learning, phase-retrieval time-frequency analysis, and sampling on t-design curves, and discussed open problems in the field.

Mathematics Subject Classification (2020): 68Txx, 42-08, 65Txx, 94Axx.

Introduction by the Organizers

The workshop Applied Harmonic Analysis and Data Processing was organized by Ingrid Daubechies, Gitta Kutyniok, and Holger Rauhut. This meeting was attended by 47 participants – thereof 18 female – from 10 countries; 46 of them participated in person and 1 participated virtually.

Data Science encompasses signal and image processing, data processing and machine learning. On the one hand it is a field of major importance for science, technology and society and on the other hand it is a very rich source of a large variety of mathematical problems. A major challenge is the ever increasing size and complexity of data and the demand for efficient computational methods for processing such data. Mathematical understanding of the underlying structures

and algorithms is highly desired. One of the key drivers for a large number of big data applications is deep learning. Despite its huge success, however, mathematical theory for major aspects of deep learning is still in its beginnings. Although highly exciting mathematical results could be shown in the recent years, many open problems remain. This means that the amount of new mathematical challenges arising from the need of data analysis and information processing is enormous, with their solution requiring fundamentally new ideas and approaches, with significant consequences in the practical applications.

This workshop was a concerted effort to bring together researchers with various backgrounds, including harmonic analysis, optimization, probability theory, approximation theory, machine learning, computer science and electrical engineering. The workshop featured 30 talks, thereof several longer overview talks. Moreover, a session of short presentations of 3 minutes took place on Monday, which we call the 3 Minutes of Fame (following Andy Warhol's concept of 15 minutes of fame). This session has meanwhile become a tradition and has proven to be an efficient vehicle to ensure that every participant had the possibility to advertise her/his research. At the same time it is very entertaining for the audience. A large part of the attendees participated, ranging from PhD students to renowned professors, contributing to the success of this session.

Let us mention a few highlights from the program:

- **Mathematical theory of deep learning.** A number of talks reported on progress – but also on intriguing open questions – on the theory of deep learning. Helmut Boelcskei talked about the relation of fuzzy logic with neural networks. Mahdi Soltanolkotabi and Carola Schönlieb reported on mathematical progress on the use of deep learning methods for the solution of inverse problems. Kathlén Kohn presented algebraic properties on linear convolutional neural networks. Rémi Gribonval showed new results on conservation laws for gradient flows, in particular for those related to learning neural networks. Sophie Langer talked about the role of statistics in deep learning. Rima Alaifari presented mathematical approaches for the use of neural networks in solving PDEs, leading to the notion of neural operators. Anna Shalova gave insights on the effect of random noise on the implicit bias of stochastic gradient descent for learning overparameterized models. Claire Boyer presented an overview on physics-informed neural networks. Noam Razin talked about implicit bias of policy gradient methods arising in reinforcement learning. Felix Voigtländer presented new results on sampling numbers of the Fourier-analytic Barron space often arising in the theoretical analysis of neural networks. Interestingly, as a main proof technique he used compressive sensing approaches. Mariia Seleznova talked about new insights as well as limitation of the kernel regime of deep neural networks.
- **Harmonic analysis approaches for signal and image processing.** Karlheinz Gröchenig presented the concept of sampling on t-design curves approximating integrals over the sphere via line integrals (modeling sensors

that move in time). He gave interesting constructions of t-design curves and presented a number of open problems. Philipp Grohs reported on stability and sampling results for phase retrieval, i.e., for reconstruction from time-frequency samples without phase information. Interestingly, this problem behaves very differently from reconstruction when the phase-information is present. Monika Dörfler talked about quantum harmonic analysis and how it can be used to augment data sets in a meaningful way.

- **Mathalchemy.** In an entertaining evening lecture, Ingrid Daubechies talked about her experience in creating an impressive artwork – called Mathalchemy – together with 23 other mathematicians and artists. It illustrates a huge number of mathematical concepts and results in a very creative installation combining many pieces of art made with a large number of techniques. This art installation is displayed at changing museums.

As a cultural activity, on Thursday evening, a concert in the great music room of MFO was organized where the group had the pleasure of listening to musical contributions at very high level by several participants ranging from Jazz to pop and classical music.

The organizers would like to take the opportunity to thank MFO for providing support and a very inspiring environment for the workshop.

Workshop: Applied Harmonic Analysis and Data Science

Table of Contents

Abstracts

Ten open problems involving matrices, randomness, graphs, and more

Afonso S. Bandeira

Below I list ten conjectures that I mentioned in my talk at Oberwolfach, and an extra open problem I thought about while there. The statements are at times stated in a slightly informal way for easiness and brevity of exposition, formal statements are available in the references. After each group of related conjectures, a short description with references is given. For the sake of brevity, the reference list focuses on recent work (that itself cites other relevant references). Some references are forthcoming, but should hopefully be available in the next few months. As always, spending time at Oberwolfach was a true pleasure, full of stimulating lectures and mathematical discussions.

Conjecture 1 (Matrix Spencer)**.** *There exists a positive universal constant C such that, for all positive integers n, and all choices of n self-adjoint $n \times n$ real matrices $A_1, \ldots, A_n$ satisfying, for all $i \in [n]$, $\|A_i\| \leq 1$ (where $\|\cdot\|$ denotes the spectral norm) the following holds*

$$\min_{\varepsilon \in \{\pm 1\}} \left\| \sum_{i=1}^{n} \varepsilon_i A_i \right\| \leq C\sqrt{n}.$$

Conjecture 2 (Group Spencer)**.** *Let G be a finite group of size n. Conjecture 1 holds in the particular case in which $A_1, \ldots, A_n$ are the $n \times n$ matrices corresponding to the regular representation of G.*

Bansal, Jiang, and Meka [7] showed Conjecture 1 for low-rank matrices and Bandeira, Kunisky, Mixon, Zeng [5] showed Conjecture 2 for simple groups.

Conjecture 3 (Kikuchi Spectral Threshold)**.** *Given r, ℓ, n positive integers satisfying $n \gg \ell \geq \frac{r}{2}$ and r even (r and ℓ will be fixed and $n \to \infty$). For each $S \subset [n]$ with $|S| = r$ let $Z_S \sim \mathcal{N}(0,1)$, and all independent. Let $\lambda \geq 0$ and let M be the $\binom{n}{\ell} \times \binom{n}{\ell}$ matrix (with rows and columns indexed by subsets $I \subset [n]$ of size ℓ) given by*

$$M(\lambda)_{I,J} = \begin{cases} \lambda + Z_{I \Delta J} & \text{if} \quad |I \Delta J| = r, \\ 0 & \text{otherwise,} \end{cases}$$

where $I \Delta J = (I \cup J) \setminus (I \cap J)$ denotes the symmetric difference.

Let $\lambda^{\natural}_{r,\ell}$ denote the threshold at which eigenvalues "pop-out" of the spectrum of $M(\lambda)$: in other words $\lambda^{\natural}_{r,\ell}$ is the real number such that, for all $\lambda > \lambda^{\natural}_{r,\ell}$, there exists $\varepsilon > 0$ such that $\mathbb{E}\lambda_{\max} M(\lambda) > (1 + \varepsilon + o(1))\mathbb{E}\lambda_{\max} M(0)$, where $o(1)$ is a term that goes to zero as $n \to \infty$.

For fixed r, we have

$$n^{\frac{r}{4}} \lambda^{\natural}_{r,\ell} \to 0$$

as $\ell \to \infty$ (note that this is after one has taken $n \to \infty$).

This conjecture was posed in the context of understanding the behavior of the Kikuchi methods for Tensor PCA in the work of Wein, El Alaoui, and Moore [11]. The threshold has been characterized for even r and $\frac{1}{2}r \leq \ell < \frac{3}{4}r$ in forthcoming work of Bandeira, Cipolloni, Schröder, and van Handel [3], which is a follow up to the matrix concentration inequalities of Bandeira, Boedihardjo, and van Handel [2].

Conjecture 4 (Tensor Non-Commutative Khintchine I)**.** *Let $r \geq 2$ an integer and $p \geq 2$ a real. Let $T_1, \ldots, T_n$ be $d^{\otimes r}$ symmetric tensors (invariant under permutation of the r indices, $r = 2$ corresponds to symmetric matrices). Given such a tensor T, we define*

$$\|T\|_{I_p} = \sup_{v \in \mathbb{R}^d : \|v\|_p = 1} \left| \left\langle v^{\otimes r}, T \right\rangle \right|,$$

for $p = 2$ this is known as the injective norm. Let $\sigma_{AW}^2 = \sum_{i=1}^n \|T_i\|_{I_p}^2$.

The following bound holds:

$$\mathbb{E} \left\| \sum_{i=1}^n g_i T_i \right\|_{I_p} \leq C_{r,p} \ \mathrm{polylog}_{r,p}(d) \ d^{\frac{1}{2} - \frac{1}{p}} \ \sigma_{AW},$$

where g_i are iid standard gaussian random variables, and $C_{r,p}$ is a constant depending on r and p, but not on d or n (and the polylogarithmic factor may also depend on r, p).

Conjecture 5 (Tensor Non-Commutative Khintchine II)**.** *A weaker version of Conjecture 4, where under the same conditions the weaker bound is conjectured*

$$\mathbb{E} \left\| \sum_{i=1}^n g_i T_i \right\|_{I_p} \leq C_{r,p} \ \mathrm{polylog}_{r,p}(d) \ d^{\frac{1}{2} - \frac{1}{\max\{p,r\}}} \ \sigma_{AW}.$$

Forthcoming work Bandeira, Gopi, Jiang, Lucca, and Rothvoss [4] establishes the weaker upper bound $\leq C_{r,p} \ \mathrm{polylog}_{r,p}(d) \ d^{\frac{1}{2} - \frac{1}{\max\{p,2r\}}} \ \sigma_{AW}$. We note also that the factor $d^{\frac{1}{2} - \frac{1}{p}}$ is necessary. This question is tightly connected to the question of proving Non-Commutative Khintchine Inequality without the use of operator theoretic tools. We refer to [4] for more on this.

Conjecture 6 (Ellipsoid Problem - Existence)**.** *Fix any $0 < \varepsilon < 1$. For d and n positive integers such that*

$$n \leq (1 - \varepsilon) \frac{d^2}{4},$$

let $x_1, \ldots, x_n$ be iid $\mathcal{N}(0, I_{d \times d})$ random vectors in $\mathbb{R}^d$. We have

$$\mathrm{Prob}\left[\exists \, \Sigma \succeq 0 \text{ such that}, \forall_{i \in [n]}, \ x_i^T \Sigma x_i = 1 \right] \to 1,$$

as $n, d \to \infty$.

Conjecture 7 (Ellipsoid Problem - Non-existence)**.** *Fix any $\varepsilon > 0$. For d and n positive integers such that*

$$n \geq (1 + \varepsilon) \frac{d^2}{4},$$

let $x_1, \ldots, x_n$ be iid $\mathcal{N}(0, I_{d\times d})$ random vectors in $\mathbb{R}^d$. We have

$$\text{Prob}\left[\exists\, \Sigma \succeq 0 \text{ such that}, \forall_{i\in[n]},\, x_i^T \Sigma x_i = 1\right] \to 0,$$

as $n, d \to \infty$.

There are known lower and upper bounds that are optimal up to a universal constant. Furthermore, Conjectures 6 and 7 are known for an approximate version of the problem, we point the reader to Maillard and Bandeira [9] and references therein.

Conjecture 8 (Globally Synchronizing Regular Graphs)**.** *We say an $n\times n$ matrix A is globally synchronizing if the only local minima of $\mathcal{E} : \mathbb{S}^{n-1} \to \mathbb{R}$, parameterized by $\theta \in [0, 2\pi[^n$ and given by*

$$\mathcal{E}(\theta) = \frac{1}{2}\sum_{i,j=1}^{n} A_{ij}\left(1 - \cos(\theta_i - \theta_j)\right),$$

are the global minima corresponding to $\theta_i = c$, $\forall_i$. A graph G is said to be globally synchronizing if its adjacency matrix is globally synchronizing.

A uniform random 3-regular graph is globally synchronizing with high probability (probability going to 1 as $n \to \infty$).

Conjecture 9 (Global Synchrony with negative edges)**.** *Given any $\varepsilon > 0$, the $n \times n$ random matrix A with zero diagonal and off-diagonal entries given by*

$$A_{ij} = \begin{cases} 1 & \text{with probability } \frac{1}{2} + \delta \\ -1 & \text{with probability } \frac{1}{2} - \delta, \end{cases}$$

with $\delta \geq (1+\varepsilon)\sqrt{\frac{\log n}{2n}}$, is globally synchronizing with high probability (see Conjecture 8).

Conjecture 10 (Density threshold for Global Synchrony)**.** *For any $\varepsilon > 0$, there exists $n > 0$ and a graph G on n nodes such that the minimum degree of G is at least $\left(\frac{3}{4} - \varepsilon\right) n$ and G is not globally synchronizing (See Conjecture 8).*

Conjectures 8 and 10 are related to understanding spontaneous synchronization of coupled oscillators under the Kuramoto model. Abdalla, Bandeira, Kassabov, Souza, Strogatz, and Townsend [1] have shown that random 600-regular graphs are globally synchronizing with high probability. Kassabov, Strogatz, and Townsend [8] showed an upper bound on the $\frac{3}{4}$ threshold in 10. Conjecture 9 is motivated by understanding performance of the Burer-Monteiro approach for inverse problems on graphs, a higher rank version of this conjecture is shown in forthcoming work of McRae, Abdalla, Bandeira, and Boumal [10] but Conjecture 9 remains elusive.

An eleventh question: During the workshop, Radu Balan proposed a beautiful question, that captured a lot of my (and my collaborators) attention for the whole workshop. I couldn't resist the opportunity to mention it here.

Open Problem 11 (The Balan-Jiang Problem). Given a self-adjoint positive semi-definite matrix $A \in \mathbb{C}^{n\times n}$, we define

$$\gamma_+(A) = \inf\left\{ \sum_{k=1}^{p} \|x_k\|_1^2 : \ x_1, \ldots, x_p \in \mathbb{C}^n, \ A = \sum_{k=1}^{p} x_k x_k^* \right\}.$$

Given $n > 0$, define

$$C_n = \sup\left\{ \frac{\gamma_+(A)}{\|A\|_{(e,1)}} : \ A \in \mathbb{C}^{n\times n}, \ A \succeq 0, A \neq 0 \right\},$$

where $\|\cdot\|_{(e,1)}$ denotes the entrywise ℓ_1 norm. What is the asymptotic growth of C_n?

During the workshop, Mixon, Steinerberger, and myself [6] showed that $C_n \geq c\sqrt{n}$, for a universal constant $c > 0$. The best known upper bound is of order n. See references within [6] for more on this question.

Acknowledgements. Thanks to Antoine Maillard, Kevin Lucca, Andrew McRae, Dustin Mixon, Stefan Steinerberger, and Mariia Seleznova for going over an earlier version of this manuscript, spotting many typos, and suggesting improvements. Last but not least, thanks to all my collaborators with whom I have spent many enjoyable hours thinking about these problems.

References

[1] P. Abdalla, A. S. Bandeira, M. Kassabov, V. Souza, S. H. Strogatz, A. Townsend, *Expander graphs are globally synchronizing*, Preprint, available on arXiv, (2023).
[2] A. S. Bandeira, M. Boedihardjo, R. van Handel *Matrix Concentration Inequalities and Free Probability*, Inventiones Mathematicae, 234 (2023), 419–487.
[3] A. S. Bandeira, G. Cipolloni, D. Schroder, R. van Handel, *Forthcoming* (2024+).
[4] A. S. Bandeira, S. Gopi, H. Jiang, T. Rothvoss, *Forthcoming* (2024+).
[5] A. S. Bandeira, D. Kunisky, D. Mixon, X. Zeng, *On the concentration of Gaussian Cayley matrices*, Preprint, available on arXiv, (2022).
[6] A. S. Bandeira, D. G. Mixon, S. Steinerberger, *A lower bound for the Balan–Jiang matrix problem*, Preprint, available on arXiv, (2024).
[7] N. Bansal, H. Jiang, R. Meka , *Resolving Matrix Spencer Conjecture Up to Poly-logarithmic Rank*, STOC (2023).
[8] M. Kassabov, S. H. Strogatz, A. Townsend, *Sufficiently dense Kuramoto networks are globally synchronizing*, Chaos **31** (2021).
[9] A. Maillard, A. S. Bandeira, *Exact threshold for approximate ellipsoid fitting of random points*, Preprint, available on arXiv, (2023).
[10] A. McRae, P. Abdalla, A. S. Bandeira, N. Boumal, *Forthcoming* (2024+).
[11] A. Wein, A. El Alaoui, C. Moore, *The Kikuchi hierarchy and tensor PCA*, FOCS (2019).

Extracting formulae in many-valued logic from deep neural networks

Helmut Bölcskei

State-of-the-art deep neural networks exhibit impressive reasoning capabilities, e.g. in mathematical tasks [1], program synthesis [2], and algorithmic reasoning [3]. This talk reports an attempt at systematically connecting neural networks with mathematical logic. Specifically, we shall be interested in reading out logical formulae from (trained) deep neural networks.

Let us first take a step back. Consider a neural network that realizes a map $f : [0,1]^n \to [0,1]$. When the input and output variables take on two possible values only, say 0 and 1, f reduces to a Boolean function and can hence be studied by means of Boolean algebra, see e.g. [4]. Boolean functions can be realized by Boolean circuits [5]. The idea of using Boolean algebra to analyze and design Boolean circuits dates back to [6, 7] and most prominently to Shannon [8]. Specifically, this correspondence works as follows. Given a Boolean circuit, one can deduce a Boolean algebraic expression that realizes the circuit's input-output relation. Conversely, for a given Boolean algebraic expression, it is possible to specify a Boolean circuit whose input-output relation equals this expression.

The main aim of the research reported in this talk is to initiate the development of a generalization of this correspondence from Boolean functions $f : \{0,1\}^n \to \{0,1\}$ to general functions $f : [0,1]^n \to [0,1]$. This immediately leads to the following two questions:

(1) What is the logical system replacing Boolean logic?
(2) What is the counterpart of Boolean circuits?

As to the first question, we show that the theory of infinite-valued Łukasiewicz logic [9] provides a suitable framework for characterizing general (nonbinary) functions $f : [0,1]^n \to [0,1]$ from a logical perspective. With slight abuse of terminology, we shall refer to infinite-valued Łukasiewicz logic as many-valued (MV) logic. Based on a fundamental result [10], which characterizes the class of truth functions in MV logic—also called McNaughton functions—as continuous piecewise linear functions with integer coefficients, we show that neural networks employing the ReLU nonlinearity $\rho(x) = \max\{0, x\}$ and integer weights[1] naturally implement statements in MV logic. This answers the second question above by identifying ReLU networks as the counterpart of Boolean circuits.

In practice, trained neural networks will, however, not exhibit integer weights, unless this is explicitly enforced in the training process. Extensions of MV logic, namely Rational Łukasiewicz logic [11] and $\mathbb{R}\mathcal{L}$ [12], have truth functions that are again continuous piecewise linear, but with rational and real-valued coefficients, respectively. Such functions are likewise naturally realized by ReLU networks, but correspondingly with rational and real-valued weights.

Besides the conceptual contribution residing in the systematic development of the connection between ReLU networks and MV logic along with its extensions, we

[1]By weights, we mean the entries of the weight matrices and bias vectors associated with the network.

also devise an algorithm for extracting logical formulae from (trained) ReLU networks with integer, rational, or real-valued weights. For pedagogical reasons and to render the presentation more accessible, we first present the entire framework for MV logic and ReLU networks with integer weights, and then provide extensions to the rational and real case. In addition, we carry out a detailed comparison between our algorithm and the only two constructive procedures for converting McNaughton functions to their associated MV logical formulae available in the literature [13, 14].

The overall philosophy of viewing ReLU networks as the circuit counterpart of MV logic and its extensions is inspired by [15, 16, 17]. Specifically, Amato et al. [15, 16] pointed out that neural networks, with the clipped ReLU (CReLU) nonlinearity $\sigma(x) = \min\{1, \max\{0, x\}\}$ and rational weights, realize truth functions in Rational Łukasiewicz logic. Di Nola et al. [17] proved that CReLU networks with real weights realize truth functions in $\mathbb{R}\mathcal{L}$ logic. The universal correspondence between ReLU networks and MV logic as well as its extensions reported here along with the algorithm for extracting logical formulae from ReLU networks appear to be new.

References

[1] G. Lample, and F. Charton. *Deep learning for symbolic mathematics*, arXiv preprint arXiv:1912.01412 (2019).

[2] S. Bubeck, V. Chandrasekaran, R. Eldan, and others. *Sparks of artificial general intelligence: Early experiments with GPT-4*, arXiv preprint arXiv:2303.12712 (2023).

[3] B. Liu, J. T. Ash, S. Goel, A. Krishnamurthy, and C. Zhang. *Transformers learn shortcuts to automata*, arXiv preprint arXiv:2210.10749 (2022).

[4] S. Skyum, and L. G. Valiant. *A complexity theory based on Boolean algebra*, Journal of the ACM (JACM) Vol. **32** (2) (1985), 484–502.

[5] R. B. Boppana, and M. Sipser. *The complexity of finite functions*, Algorithms and Complexity (1990), 757–804.

[6] A. Nakashima. *The theory of relay circuit composition*, The Journal of the Institute of Telegraph and Telephone Engineers of Japan Vol. **38** (7) (1935), 461–489.

[7] V. I. Shestakov. *Some Mathematical Methods for the Construction and Simplification of Two-Terminal Electrical Networks of Class A*, PhD thesis, The Lomonosov State University, Moscow, Russia (1938).

[8] C. E. Shannon. *A symbolic analysis of relay and switching circuits*, Electrical Engineering Vol. **57** (12) (1938), 713–723.

[9] R. L. Cignoli, I. M. d'Ottaviano, and D. Mundici. *Algebraic Foundations of Many-Valued Reasoning*, Vol. **7** (2000).

[10] R. McNaughton, *A theorem about infinite-valued sentential logic*, The Journal of Symbolic Logic Vol. **16** (1) (1951), 1–13.

[11] B. Gerla, *Rational Łukasiewicz logic and DMV-algebras*, Neural Network World Vol. **6** (6) (2001).

[12] A. Di Nola, and I. Leuştean, *Łukasiewicz logic and Riesz spaces*, Soft Computing Vol. **18** (12) (2014), 2349–2363.

[13] D. Mundici, *A constructive proof of McNaughton's theorem in infinite-valued logic*, The Journal of Symbolic Logic Vol. **59** (2) (1994), 596–602.

[14] S. Aguzzoli, *Geometrical and Proof Theoretical Issues in Łukasiewicz Propositional Logics*, PhD thesis, University of Siena, Italy (1998).

[15] P. Amato, A. Di Nola, and B. Gerla, *Neural networks and rational Łukasiewicz logic*, Annual Meeting of the North American Fuzzy Information Processing Society Proceedings (2002), 506–510.
[16] P. Amato, A. Di Nola, and B. Gerla, *Neural networks and rational McNaughton Functions*, Journal of Multiple-Valued Logic and Soft Computing Vol. **11** (1-2) (2005), 95–110.
[17] A. Di Nola, B. Gerla, and I. Leuştean, *Adding real coefficients to Łukasiewicz logic: An application to neural networks*, Fuzzy Logic and Applications: 10th International Workshop (2013), 77–85.

Deep learning for inverse problems: faster reconstruction, new architectures, and dynamic feature learning

Mahdi Soltanolkotabi

This talk focuses on three interconnected themes. In the first part of the talk I discussed our recent work on flash diffusion in [2]. Inverse problems arise in a multitude of applications, where the goal is to recover a clean signal from noisy and possibly (non)linear observations. The difficulty of a reconstruction problem depends on multiple factors, such as the structure of the ground truth signal, the severity of the degradation, the implicit bias of the reconstruction model and the complex interactions between the above factors. This results in natural sample-by-sample variation in the difficulty of a reconstruction task, which is often overlooked by contemporary techniques. Recently, diffusion-based inverse problem solvers have established new state-of-the-art in various reconstruction tasks. Our key observation in this paper is that most existing solvers lack the ability to adapt their compute power to the difficulty of the reconstruction task, resulting in long inference times, subpar performance and wasteful resource allocation. We propose a novel method that we call severity encoding, to estimate the degradation severity of noisy, degraded signals in the latent space of an autoencoder. We show that the estimated severity has strong correlation with the true corruption level and can give useful hints at the difficulty of reconstruction problems on a sample-by-sample basis. Furthermore, we propose a reconstruction method based on latent diffusion models that leverages the predicted degradation severities to fine-tune the reverse diffusion sampling trajectory and thus achieve sample-adaptive inference times. We perform numerical experiments on both linear and nonlinear inverse problems and demonstrate that our technique achieves performance comparable to state-of-the-art diffusion-based techniques, with significant improvements in computational efficiency.

In the second part I talked about our recent work on designing new architectures for accelerated MRI reconstruction [1]. In accelerated MRI reconstruction, the anatomy of a patient is recovered from a set of under-sampled and noisy measurements. Deep learning approaches have been proven to be successful in solving this ill-posed inverse problem and are capable of producing very high quality reconstructions. However, current architectures heavily rely on convolutions, that are content-independent and have difficulties modeling long-range dependencies in images. Recently, Transformers, the workhorse of contemporary natural language

processing, have emerged as powerful building blocks for a multitude of vision tasks. These models split input images into non-overlapping patches, embed the patches into lower-dimensional tokens and utilize a self-attention mechanism that does not suffer from the aforementioned weaknesses of convolutional architectures. However, Transformers incur extremely high compute and memory cost when 1) the input image resolution is high and 2) when the image needs to be split into a large number of patches to preserve fine detail information, both of which are typical in low-level vision problems such as MRI reconstruction, having a compounding effect. To tackle these challenges, we propose HUMUS-Net, a hybrid architecture that combines the beneficial implicit bias and efficiency of convolutions with the power of Transformer blocks in an unrolled and multi-scale network. HUMUS-Net extracts high-resolution features via convolutional blocks and refines low-resolution features via a novel Transformer-based multi-scale feature extractor. Features from both levels are then synthesized into a high-resolution output reconstruction. Our network establishes new state of the art on the largest publicly available MRI dataset, the fastMRI dataset. We further demonstrate the performance of HUMUS-Net on two other popular MRI datasets and perform fine-grained ablation studies to validate our design.

Finally, in the last part of the talk I discussed our recent work on feature learning via gradient descent in [3]. We consider the problem of learning polynomials with low-dimensional latent representation of the form $f^*(x) = g(Ux)$, where U maps from d to r dimensions with $d \gg r$. When g is a degree g polynomial, existing analysis using equivalence to kernel methods require $n \asymp d^p$ samples. Our primary result is that gradient descent can learn the representation $\text{span}(U)$, and then perform a kernel method restricted to the span. This results in a sample complexity of $n \asymp d^2 r + dr^p$. Finally, we show that the assumptions in our algorithm is necessary by showing a correlational statistical query lower bound of $n \asymp d^{O(p)}$ when the assumptions are violated.

References

[1] F. Zalan, B. Tinaz, M. Soltanolkotabi, *Humus-net: Hybrid unrolled multi-scale network architecture for accelerated mri reconstruction*, Advances in Neural Information Processing Systems 35 (2022), 25306-25319.

[2] F. Zalan, B. Tinaz, M. Soltanolkotabi, *Adapt and Diffuse: Sample-adaptive Reconstruction via Latent Diffusion Models*, International Conference on Machine Learning Research (ICML) (2024).

[3] D. Alexandru, J. Lee, M. Soltanolkotabi, *Neural networks can learn representations with gradient descent*, Conference on Learning Theory, PMLR (2022).

t-design curves and mobile sampling on the sphere

Karlheinz Gröchenig
(joint work with Martin Ehler, Clemens Karner)

A finite set $X_t \subseteq \mathbb{S}^d$ is a t-design (or X_t consists of t-design points), if for every algebraic polynomial f in $d+1$ variables of (total) degree t one has

$$\frac{1}{|X_t|} \sum_{x \in X_t} f(x) = \int_{\mathbb{S}^d} f \,. \tag{1}$$

Using curves instead of points, one can define analogously the notion of a t-design curve on the sphere $\mathbb{S}^d$ as follows [2]. A piecewise smooth curve $\gamma : [0,1] \to \mathbb{S}^d$ with at most finitely many self intersections and and with arc length $\ell(\gamma)$ is called a t-design curve in $\mathbb{S}^d$, if the line integral integrates exactly all algebraic polynomials in $d+1$ variables of degree t, i.e.,

$$\frac{1}{\ell(\gamma)} \int_\gamma f = \int_{\mathbb{S}^d} f \,. \tag{2}$$

Here $\int_\gamma f = \int_0^1 f(\gamma(s))|\gamma'(s)|\, ds$ and the arc length is $\ell(\gamma) = \int_\gamma 1$.

The definition is motivated by Martin Vetterli's idea of mobile sampling [4]. Instead of using many fixed sensors to measure a physical field one might want to use a single sensor and move it around, in other words, instead of sampling at many points one should sample along a curve. Thus t-design curves can be seen as a version of mobile sampling on the sphere.

Very little is known about t-design curves on the sphere. There are some ad hoc constructions of t-design curves for $t = 1, 2, 3$, some necessary conditions on the required length of t-design curves, and some existence results. We report on our work [2].

A necessary condition is the following.

Theorem 1. *Assume that a piecewise smooth, closed curve $\gamma : [0,1] \to \mathbb{S}^d$ satisfies*

$$\frac{1}{\ell(\gamma)} \int_\gamma f = \int_{\mathbb{S}^d} f \qquad \text{for all } f \in \Pi_t \,.$$

Then its length is bounded from below by

$$\ell(\gamma) \geq C_d t^{d-1}$$

with some constant $C_d > 0$ that may depend on the dimension d but is independent of t and γ.

For comparison, the number of t-design points is bounded below by $C'_d t^d$ for some dimensional constant.

It was a longstanding open problem to prove the existence of t-design points achieving this lower bound. Attempts in the 1990 showed the existence of t-design points with $\mathcal{O}(t^{d^3})$ and then $\mathcal{O}(t^{d^2})$ points. The breakthrough result is due to A. Bondarenko, D. Radchenko, and M. Viazhovska [1] who proved the existence of t-design points of the conjectured order $\mathcal{O}(t^d)$.

As for existence of t-design curves, we were able to use the result of Bondarenko, Radchenko, and Viazhovska. For the 2-sphere our result is asymptotically optimal, for higher-dimensional spheres the result corresponds to versions for t-design points before [3].

Theorem 2. *In $\mathbb{S}^2$ there exists a sequence of t-design curves $(\gamma_t)_{t\in\mathbb{N}}$ with length $\ell(\gamma_t) \asymp t$.*

In view of Theorem 1 this result is asymptotically optimal. The result for higher-dimensional spheres is obtained by induction on the dimension, and in this process one looses the asymptotic optimality.

Theorem 3. *In $\mathbb{S}^d$ for $d \geq 3$ there exists a sequence of t-design curves $(\gamma_t)_{t\in\mathbb{N}}$, such that $\ell(\gamma_t) \lesssim t^{d(d-1)/2}$.*

Both results are based on the existence of asymptotically optimal t-design points from [1]. The technical part involves some harmonic analysis and geometry on the sphere, and in particular a lemma of Samko [5] that connects point evaluations to line integrals along circles. The final step is the construction of a *closed* curve and requires some elementary graph theory, such as the existence of Euler paths.

A natural idea would be to start with a set of t-design points and to connect them along geodesic arcs. By choosing a suitable sequence of points, one could construct a curve of the desired length $\mathcal{O}(t^{d-1})$ and then hope that miraculously this curve also yields exact integration, i.e., is a t-design curve. Unfortunately, this idea fails, as already the examples for $t = 2$ and $t = 3$ show. In fact, this was our first idea, and this is the audience's guess in every talk and every discussion. However, this procedure yields curves satisfying Marcinkiewicz-Zygmund Inequalities.

Theorem 4. *Let $X_t \subseteq \mathbb{S}^d$ be a sequence of t-design points as constructed in* [1]. *Then there exists a closed, piecewise smooth curve γ_t with vertices at X_t of length $\ell(\gamma_t) \asymp t^{d-1}$, such that*

$$A \int_{\mathbb{S}^d} |f|^2 \leq \frac{1}{\ell(\gamma_t)} \int_{\gamma_t} |f|^2 \leq B \int_{\mathbb{S}^d} |f|^2 \tag{3}$$

for all polynomials of degree t.

This result is still sufficient for many purposes, e.g., it implies a formula for exact integration and yields sufficiently dense curves.

The investigation and construction of t-design curves is only at the beginning and, for the authors, a big, new research project. There are many open problems and many variations to be looked at.

References

[1] A. Bondarenko, D. Radchenko, and M. Viazovska, *Optimal asymptotic bounds for spherical designs*, Ann. of Math. (2) (2013), 178(2):443–452 .

[2] M. Ehler and K. Gröchenig, *t-design curves and mobile sampling on the sphere*, Forum Math. Sigma (2023), 11:Paper No. e105, 25, .

[3] J. Korevaar and J. L. H. Meyers, *Spherical Faraday cage for the case of equal point charges and Chebyshev-type quadrature on the sphere*, Integral Transform. Spec. Funct. (1993), 1(2):105–117.
[4] J. Unnikrishnan and M. Vetterli, *Sampling high-dimensional bandlimited fields on low-dimensional manifolds*, IEEE Trans. Inform. Theory (2013), 59(4):2103–2127.
[5] S. G. Samko, *Generalized Riesz potentials and hypersingular integrals with homogeneous characteristics; their symbols and inversion*, Trudy Matematicheskogo Instituta imeni VA Steklova (1980), 156:157-222.

Convolutions, groups, polynomials in neural networks

Kathlén Kohn
(joint work with Guido Montúfar, Anna-Laura Sattelberger, Vahid Shahverdi, Matthew Trager)

Feedforward neural networks are parametrized families of functions. For a fixed network architecture, its weights and biases are the trainable parameters $\theta \in \mathbb{R}^N$. Each layer of the network is a function $f_{i,\theta}$ that depends on the parameters θ. A network with L layers composes L such functions via the *network parametrization map*

$$\mu : \mathbb{R}^N \longrightarrow \mathcal{M}, \quad \theta \longmapsto f_{L,\theta} \circ \ldots \circ f_{2,\theta} \circ f_{1,\theta}. \tag{1}$$

The space of functions $\mathcal{M} = \mathrm{im}(\mu)$ that the network architecture parametrizes is often called the *neuromanifold* of the network, although it is not a smooth manifold (it essentially always has singularities).

Given training data $\mathcal{D}$, supervised learning aims to minimize the loss

$$\mathcal{L}_{\mathcal{D}} : \mathbb{R}^N \xrightarrow{\mu} \mathcal{M} \xrightarrow{\ell_{\mathcal{D}}} \mathbb{R}. \tag{2}$$

Typically, a version of gradient descent is used to find parameters $\theta \in \mathbb{R}^N$ such that $\mu(\theta) \in \mathcal{M}$ is a "best" function approximating the training data $\mathcal{D}$.

This optimization problem has a deeply geometric nature since the loss $\mathcal{L}_{\mathcal{D}}$ is a composition of the network parametrization map μ and the loss $\ell_{\mathcal{D}}$ on function space. This raises the questions:

- How does the network architecture affect the geometry of the neuromanifold $\mathcal{M}$?
- How does the geometry of $\mathcal{M}$ impact the training of the network?

1. Impact of convolutional architecture

We begin by investigating how the neuromanifold and the training of the network are affected by changing a linear network from a dense architecture to a convolutional architecture.

In a *linear dense network*, the functions in (1) are arbitrary linear maps $f_{i,\theta} : \mathbb{R}^{k_{i-1}} \to \mathbb{R}^{k_i}$. The parameters $\theta \in \mathbb{R}^{k_L \times k_{L-1}} \times \ldots \times \mathbb{R}^{k_1 \times k_0}$ are the entries of the matrices representing $f_{L,\theta}, \ldots, f_{1,\theta}$. The neuromanifold is $\mathcal{M} = \{W \in \mathbb{R}^{k_L \times k_0} \mid \mathrm{rank}(W) \leq \min(k_0, \ldots, k_L)\}$. This is an *algebraic variety*.

Next, we consider *linear convolutional networks (LCNs)* with one channel and convolutions on one-dimensional signals. In this case, each function $f_{i,\theta}$ is a linear map, given by a *Toeplitz matrix of stride s* of the form

$$T_{w,s} = \begin{bmatrix} w_0 & \cdots & w_s & \cdots & w_{k-1} & & & & \\ & & w_0 & & \cdots & & w_{k-1} & & \\ & & & \ddots & & & & \ddots & \\ & & & & & w_0 & & \cdots & w_{k-1} \end{bmatrix}.$$

Here, $w \in \mathbb{R}^k$ is the *filter* of the convolution; all other entries in $T_{w,s}$ are 0. The composition of convolutions $f_{L,\theta} \circ \ldots \circ f_{1,\theta}$ is also a convolution: it has stride $s_1 \cdots s_L$ and its filter can be computed via sparse polynomial factorization as follows. Denoting by $w^{(i)} \in \mathbb{R}^{k_i}$ the filter of the convolution $f_{i,\theta}$ in the i-th layer and by s_i its stride, we identify the convolution with the polynomial

$$P_i(x) := w_0^{(i)} + w_1^{(i)} x^{t_i} + \ldots + w_{k_i-1}^{(i)} (x^{t_i})^{k_i-1}, \quad \text{where } t_i := s_1 \cdot s_2 \cdots s_{i-1}.$$

Under this identification, the filter of the end-to-end convolution $f_{L,\theta} \circ \ldots \circ f_{1,\theta}$ is the coefficient vector of the polynomial $P_L(x) \cdots P_1(x)$. The identification of convolutions on higher-dimensional signals with *multivariate* polynomials works analogously [1, Sections 4.2-4.3]. For the case of one-dimensional convolutions, we in particular see that the filter of the end-to-end convolution has size $k := k_1 + \sum_{i=2}^{L} (k_i - 1)t_i$. Using elementary algebraic geometry, we can understand the geometry of the set of all end-to-end filters:

Theorem 1. [2, Theorem 2.4] *The neuromanifold $\mathcal{M}$ of an LCN is a semi-algebraic Euclidean-closed subset of $\mathbb{R}^k$ of dimension $k_1 + \ldots + k_L - (L-1)$.*

In particular, we see that the neuromanifold $\mathcal{M}$ is a full-dimensional subset of its ambient vector space $\mathbb{R}^k$ if and only if all strides s_i are 1 (except for the last stride s_L, which does not affect the geometry of $\mathcal{M}$ at all). We compare the properties of $\mathcal{M}$ for dense and convolutional networks in the following table:

The table says that the neuromanifolds of linear dense networks and LCNs of stride one have degenerate points, which are singularities or boundary points, respectively. Moreover, with positive probability on the training data, some critical points of the loss correspond to such degenerate points in function space. On the contrary, LCNs with all strides larger than 1 have both types of degenerate points on their neuromanifold, but those correspond almost never to critical points when we train with the squared-error loss. This can be formulized as follows:

Theorem 2. [2, Theorem 2.12] *Consider an LCN with all strides larger than* 1 *and the squared-error loss $\ell_{\mathcal{D}}$. Let $d \geq k$. For almost all data $\mathcal{D} \in (\mathbb{R}^{k_0} \times \mathbb{R}^{k_L})^d$, every critical point θ of $\mathcal{L}_{\mathcal{D}}$ satisfies:*

- *either $\mu(\theta) = 0$, or*
- *$\mu(\theta)$ is a smooth point in the relative interior of $\mathcal{M}$ and θ is a regular point of μ.*

Hence, the network parametrization does not introduce any additional critical points (except for the zero functions) and we can equivalently study the critical points of the loss in either parameter or function space. This can be used to

	linear dense	LCN $\forall i : s_i = 1$	LCN $\forall i : s_i > 1$
$\mathcal{M}$	algebraic variety	semialgebraic & Euclidean closed	
		full-dimensional	low-dimensional
$\partial\mathcal{M}$	$\emptyset$	non-empty	non-empty
$\mathrm{Sing}(\mathcal{M}^\circ)$	non-empty	$\emptyset$	non-empty
$\mu(\mathrm{Crit}(\mathcal{L}_\mathcal{D}))$	often in $\mathrm{Sing}(\mathcal{M})$	often in $\partial\mathcal{M}$	almost never in $\mathrm{Sing}(\mathcal{M}^\circ)$ or $\partial\mathcal{M}$
critical points spurious?	often	often	almost never

TABLE 1. $\partial\mathcal{M}$ is the Euclidean relative boundary of $\mathcal{M}$. $\mathrm{Sing}(\mathcal{M}^\circ)$ refers to the singular points of $\mathcal{M} \setminus \partial\mathcal{M}$. The second to last row concerns the image of the critical points of the loss in (2) under the map μ in (1). A critical point θ of the loss $\mathcal{L}_\mathcal{D}$ is called *spurious* if $\mu(\theta)$ is not a critical point of $\ell_\mathcal{D}$. See [2].

count all critical points [4]. For a detailed study of the singularities and boundary points and the generalization of Theorem 2 to architectures where some strides are allowed to be 1, see [2].

2. Group-equivariant architectures

Imagine you want to design a linear autoencoder $f : \mathbb{R}^k \to \mathbb{R}^r \to \mathbb{R}^k$ that is equivariant under the action of a group G, i.e., such that $f(g \cdot x) = g \cdot f(x)$ for all $g \in G$ and $x \in \mathbb{R}^k$. Is it possible to find an architecture such that its neuromanifold is the whole set of all possible linear equivariant functions

$$\{f : \mathbb{R}^k \to \mathbb{R}^k \mid \text{linear}, G\text{-equivariant}, \mathrm{rank}(f) \leq r\}\,? \tag{3}$$

We show that the answer is typically *No*. More concretely, the variety (3) typically has many irreducible components. Since any neural network can parametrize at most one of those components, there is no neural network who can parametrize all of (3). This means that any architecture design would implicitly pick one of the components, and completely disregard the others. In practice, it is also not possible to design one architecture for each component and then train all those networks in parallel, because the number of components will typically be gigantic. For instance, for autoencoders trained on MNIST with $k = 784$ and $r = 99$ that are equivariant under horizontal shifts by one pixel, the number of components is

$$72,425,986,088,826.$$

A formula that computes the number of components for arbitrary k, r and for cyclic permutation groups is provided in [3]. A central question remains: **Which of these many components is best?**

3. Neural-network computable?

Can one develop a theory of which problems can be solved well by training neural networks and which one cannot, similarly to P vs. NP or quantum complexity theory? **For instance, is solving systems of polynomial equations hard for neural networks?** The classical computer vision problem of reconstructing 3D scenes from images taken by unknown cameras is essentially equivalent to solving (certain) polynomial equation systems, and so far pure machine learning solvers do not work as well classical algorithms at the core.

References

[1] K. Kohn, T. Merkh, G. Montúfar, M. Trager, *Geometry of Linear Convolutional Networks*, SIAM Journal on Applied Algebra and Geometry **6** (2022), 368–406.
[2] K. Kohn, G. Montúfar, V. Shahverdi, M. Trager, *Function Space and Critical Points of Linear Convolutional Networks*, accepted at SIAM Journal on Applied Algebra and Geometry, arXiv:2304.05752.
[3] K. Kohn, A.-L. Sattelberger, V. Shahverdi, *Geometry of Linear Neural Networks: Equivariance and Invariance under Permutation Groups*, arXiv:2309.13736.
[4] V. Shahverdi, *Algebraic Complexity and Neurovariety of Linear Convolutional Networks*, arXiv:2401.16613.

Conservation laws for gradient flows

Rémi Gribonval
(joint work with Sibylle Marcotte, Gabriel Peyré)

We rigorously expose the basic definitions and properties of "conservation laws", that define quantities conserved during gradient flows of a given machine learning model, such as a ReLU network, with any training data and any loss. We explain how to find the maximal number of independent conservation laws via Lie algebra computations and provide algorithms algorithms to: a) compute a family of polynomial laws; b) compute the number of (not necessarily polynomial) conservation laws. We obtain that on a number of architecture there are no more laws than the known ones [1], and we identify new laws for certain flows with momentum and/or non-Euclidean geometries [2].

References

[1] S. Marcotte, R. Gribonval, G. Peyré, *Abide by the Law and Follow the Flow: Conservation Laws for Gradient Flows*, 37th Conference on Neural Information Processing Systems (NeurIPS) (2023).
[2] S. Marcotte, R. Gribonval, G. Peyré, *Keep the Momentum: Conservation Laws beyond Euclidean Gradient Flows*, Forty-first International Conference on Machine Learning (ICML) (2024).

Learning to predict ground state properties of gapped Hamiltonians

Richard Küng

(joint work with Hsin-Yuan (Robert) Huang, Giacomo Torlai, Victor Albert and John Preskill)

Classical machine learning (ML) provides a potentially powerful approach to solving challenging quantum many-body problems in physics and chemistry. However, the advantages of ML over traditional methods have not been firmly established. In this talk, I sketch a rigorous proof that underscores how classical ML algorithms can efficiently predict ground-state properties of gapped Hamiltonians after learning from other Hamiltonians in the same quantum phase of matter. By contrast, under a widely accepted conjecture, classical algorithms that do not learn from data cannot achieve the same guarantee.

Our proof technique combines five main ingredients:

(1) lifting the originally quadratic 2^n-dimensional vector problem to a linear problem on Hermitian 2^n times 2^n matrices;
(2) employing importance sampling and tensor calculus to efficiently store and process Hermitian 2^n times 2^n matrices, dubbed a classical shadow;
(3) linearizing the action of a deep neural network by using recent insights about the neural tangent kernel;
(4) approximating a slowly-varying and bounded function with independent Fourier sampling (non-asymptotic);
(5) justifying the features in (4) by using tensor product expansions, as well as matrix perturbation theory. I will try to convey how these main proof ingredients arise and play nicely together in the setup of learning ground state properties of gapped local Hamiltonians. Numerical experiments that address the anti-ferromagnetic Heisenberg model and Rydberg atom systems lend further credence to our findings.

References

[1] H. Y. Huang, R. Kueng, G. Torlai, V. V. Albert, J. Preskill, *Provably efficient machine learning for quantum manybody problems*, Science (2021).

Understanding deep learning: Beyond feedforward neural networks in nonparametric regression

Sophie Langer

(joint work with Johannes Schmidt-Hieber, Harro Walk, Alina Braun, Michael Kohler)

In recent years, substantial efforts have been dedicated to gaining a deeper understanding of deep learning methods, primarily focusing on vanilla feedforward neural networks (FNNs). These networks characterized by L hidden layers and r

neurons per layer, can be described by a parametrized family of functions denoted as $\mathbf{x} \to f(\mathbf{x};\theta)$:

$$f(\mathbf{x};\theta) := W_L \sigma_L W_{L-1} \sigma_{L-1} \cdots W_1 \sigma_1 W_0 \mathbf{x}.$$

Here $\theta = (W_1, \ldots, W_L)$ with each W_i representing a $d_{i+1} \times d_i$ weight matrix and $\sigma_i : \mathbb{R}^{d_i} \to R^{d_i}$ being a fixed so-called activation function, typically the ReLU function $\sigma(x) = \max\{x, 0\}$. Given some training data $\{(x_1, y_1), \ldots, (x_n, y_n)\}$, the computation of the parameter θ is usually achieved by approximately minimizing the empirical risk

$$\hat{R}(f) = \frac{1}{n} \sum_{i=1}^{n} \ell\left(y_i, f(\mathbf{x}_i;\theta)\right)$$

for some loss function $\ell : Y \times Y \to \mathbb{R}$ using variants of gradient descent. Existing theoretical results on FNNs include the analysis of their expressive power (e.g., [1, 2, 3]), their generalization abilities on unknown new data sets (e.g., [4, 5]) and investigations in (stochastic) gradient descent for their training (e.g., [6, 7]).

While these results focus on specific aspects of the methodology, they are lacking a comprehensive end-to-end analysis of all aspects together. This gap, in turn, limits a holistic theoretical characterization of the entire procedure. In this context, the statistical perspective is a very promising approach. By interpreting deep learning as a nonlinear or nonparametric generalization of existing statistical models, networks can be analysed as estimators for statistical prediction problem such as nonparametric regression or classification.

In particular, $\hat{f}$ is defined as a neural network estimator trained by variants of gradient descent on n independent and identically distributed random variables $\mathcal{D}_n = \{(X_1, Y_1), \ldots, (X_n, Y_n)\}$ over a class of neural networks $\mathcal{F}$ with a specific architecture. The performance of the estimator $\hat{f}$ is then theoretically evaluated by analysing its (theoretical) *risk*

$$R(\hat{f}) = \mathbf{E}\left\{\ell\left(Y, \hat{f}(\mathbf{X})\right)\right\}$$

relatively to the risk of the minimizer $f^* \in \arg\min_f R(f)$ over all measurable functions and for a new random input-output pair $(\mathbf{X}, Y)$ following the same distribution as $\mathcal{D}_n$. By defining $\hat{f}_{erm} \in \arg\min_{f\in\mathcal{F}} \hat{R}(f)$ as an empirical risk minimizer over the class $\mathcal{F}$ and $R_{\mathcal{F}} = \inf_{f\in\mathcal{F}} R(f)$ as the minimal risk over $\mathcal{F}$, the error can be bounded by three terms:

$$R(\hat{f}) - R(f^*) \leq \underbrace{2 \sup_{f\in\mathcal{F}} R(f) - \hat{R}(f)}_{\text{generalization}} + \underbrace{\hat{R}(\hat{f}) - \hat{R}(\hat{f}_{erm})}_{\text{optimization}} + \underbrace{R_{\mathcal{F}} - R(f^*)}_{\text{approximation}}.$$

Here, the *approximation error* computes the richness of the function class $\mathcal{F}$, determining if the problem can be effectively computed by functions within $\mathcal{F}$. The *optimization error* measures the performance of the optimization algorithm, specifically a variant of gradient descent, and the *generalization error* evaluates

how well the empirical risk minimizer performs relative to the true risk minimizer in $\mathcal{F}$ by analysing the statistical complexity of the function class.

While the statistical perspective enables to address all three aspects together, existing results ([5, 8, 9]) often simplify their analysis by (i) restricting their function class to simple FNNs and (ii) presuming an optimal optimization, wherein $\hat{f} = f_{erm}$. While this limitation played a pivotal role in providing preliminary insights into deep learning methodologies, there is now an undeniable necessity to challenge these assumptions in order to derive theoretical results more aligned with the practical methodologies in use.

In this talk, we propose ideas for extending the existing theory, focusing on *two* aspects:

(i) The analysis of convolutional neural networks (CNNs) in image classification
(ii) The analysis of neural networks trained by gradient descent.

For (i) we present a new statistical framework for analysing image classification theoretically. Conventional nonparametric regression is often used to analyze estimators in statistical contexts [10], which assumes a functional relation between input and output, i.e., $Y = f(X) + \epsilon$, up to some additive measurement noise ϵ. However, this model becomes questionable for image classification due to the high variability of images within a single class, leading to high prediction errors due to the high noise values in the data.

To address this challenge, we propose a new image classification model where images are viewed as highly structured objects, and variations within a class are attributed to different geometric deformations. The classification rule should, therefore, remain invariant to these deformations. While CNNs are inherently translation-invariant, their learning of other geometric deformations needs theoretical exploration. In this report, we provide an initial theoretical foundation for CNN-based image recognition, emphasizing their superior performance in handling a variety of geometric deformations.

To (ii) analyse neural networks trained by gradient descent, we start the analysis with shallow neural networks and regression functions lying in the so-called Barron class - functions with a finite first moment of their Fourier transform. In this context, Andrew Barron showed convergence guarantees for an empirical risk minimizer based on shallow networks ([11]). Our results extend his analysis to networks trained by gradient descent. An important aspect in our analysis is the proper initilization of the weights in the hidden layer of the network. Using a sigmoidal activation function, we introduce an initialization that achieves good convergence guarantees for the corresponding shallow network estimator. We further show that these weights change only slightly during gradient descent, thus providing convergence guarantees for the estimator learned by gradient descent.

In summary, this report aims to open new avenues for understanding deep learning theoretically, presenting initial ideas for potential research directions.

References

[1] G. Montufar, *Universal Approximation Depth and Errors of Narrow Belief Networks with Discrete Units*, Neural Computation **26** (2014), 1386 - 1407.
[2] R. Eldan and O. Shamir, *The Power of Depth for Feedforward Neural Networks*, COLT (2016), 907–940.
[3] D. Yarotsky, *Error bounds for approximations with deep ReLU networks*, arxiv: 1610.01145 (2016).
[4] P. L. Bartlett and N. Harvey and C. Liaw and A. Mehrabian, *Nearly-tight VC-dimension bounds for piecewise linear neural networks*, arxiv: 1703.02930 (2017).
[5] J. Schmidt-Hieber, *Nonparametric regression using deep neural networks with ReLU activation function*, Ann. Statist. **48** (4) (2020), 1875-1897.
[6] S. Pesme, N. Flammerion, *Saddle-to-saddle dynamics in diagnal linear networks*, 37th Conference on Neural Information Processing (2023).
[7] L. Chizat and F. Bach, *Implicit bias of gradient descent for wide two-layer neural networks trained with logistic loss*, 33th Conference on Learning Theory (2020).
[8] T. Suzuki, *Adaptivity of deep ReLU networks for learning Besov and mixed Besov spaces: optimal rate and curse of dimensionality*, ICLR (2019).
[9] M. Kohler, S. Langer, *On the rate of convergence of fully connected deep neural network regression estimates*, Annals of Statistics (2021).
[10] L. Györfi, M. Kohler, A. Krzyżak, H. Walk, *A Distribution-Free Theory of Nonparametric Regression*, Springer Series in Statistics (2002).
[11] A. R. Barron, *Approximation and Estimation Bounds for Artificial Neural Networks*, Machine Learning **14** (1) (1995), 115–133.

Representation equivalent Neural Operators

Rima Alaifari

(joint work with Francesca Bartolucci, Emmanuel de Bézenac, Bogdan Raonić, Roberto Molinaro, Siddhartha Mishra)

Introduction. Until recently, data-driven methods for learning maps between infinite-dimensional spaces (especially relevant in the context of solving PDEs) have followed the strategy to simply discretize both input and output spaces and learn finite-dimensional maps. Such approaches have a major drawback: the map learned in the training procedure (for a fixed discretization) may not generalize to different discretizations after training. To circumvent this shortcoming, *neural operators* [2] have recently been introduced as a concept of a learned map between *infinite-dimensional* spaces consisting of layers of linear integral operators followed by non-linear activation functions. They are defined through *discretization invariance:* (i) any discretization of the input function is accepted; (ii) the output function can be evaluated at any discretization; (iii) the discretized maps tend to an operator mapping between infinite-dimensions as the discretization becomes finer.

Imposing discretization invariance on the neural operator, however, does not enforce any link between the discretized map for a *fixed discretization* and the neural operator (as a map between infinite-dimensional spaces). It is rather an *asymptotic property* and numerical experiments show that significant discrepancies (through

aliasing errors) can occur when such architectures are tested at resolutions that differ from the training resolution.

To overcome the limitations of mere discretization invariance, we develop a novel framework of *Representation equivalent Neural Operators (ReNOs)* [1], that rather guarantees a *continuous-discrete equivalence.* Our construction is based on utilizing *frame theory* and generalizing aliasing errors to operators. This way, unique and stable reconstruction of the infinite-dimensional operator from its discretizations can be guaranteed. By connecting *all* discretizations to the same operator between infinite-dimensions, any two discretizations are also tied together uniquely and stably. Numerical experiments further highlight that there is no trade-off between the expressivity of the network and its ReNO property. As a novel ReNO architecture we propose *Convolutional Neural Operators* [3], that achieve state-of-the-art results on a large set of benchmark PDEs.

Frame Theory and Aliasing. To facilitate the transition from continuous to discrete representations we employ frame theory.

A countable sequence $\{f_i\}_{i\in I}$ in a separable Hilbert space $\mathcal{H}$ is a frame if there exist constants $A, B > 0$ such that:

$$A\|f\|^2 \leq \sum_{i\in I} |\langle f, f_i\rangle|^2 \leq B\|f\|^2, \quad \forall f \in \mathcal{H}.$$

Given such a frame, functions in $\mathcal{H}$ can be uniquely and stably reconstructed from their frame coefficients using the synthesis operator T and its adjoint T^*. The reconstruction formula is:

$$f = TT^\dagger f = \sum_{i\in I} \langle f, S^{-1}f_i\rangle f_i = \sum_{i\in I} \langle f, f_i\rangle S^{-1}f_i,$$

where $S = TT^*$, the frame operator, is invertible and its inverse S^{-1} is bounded, ensuring stability. The pseudo-inverse of the synthesis operator is given by $T^\dagger f = \{\langle f, S^{-1}f_i\rangle\}_{i\in I}$. It is useful to consider discrete representations of *subspaces* of $\mathcal{H}$, in order to quantify errors. If $\{f_i\}_{i\in I}$ is a frame of a subspace $V \subseteq \mathcal{H}$, it is said to be a *frame sequence* of $\mathcal{H}$. Aliasing occurs when a function cannot be perfectly reconstructed from its coefficients, leading to errors:

$$\epsilon(f) = f - P_V f, \quad \|\epsilon(f)\|_2 = \|f - P_V f\|_2,$$

where P_V is the orthogonal projection onto V. In other words, aliasing errors occur when a function $f \in \mathcal{H}$ with $f \notin V$ is represented by $\{\langle f, f_i\rangle\}_{i\in I}$, where $\{f_i\}_{i\in I}$ is a frame for $V \subset \mathcal{H}$, but not for $\mathcal{H}$.

Alias-Free Framework for Operator Learning. Next, we extend the concept of aliasing from functions to operators. Let U be an operator between separable Hilbert spaces $\mathcal{H}$ and $\mathcal{K}$, and let $\Psi = \{\psi_i\}_{i\in I}$ and $\Phi = \{\varphi_j\}_{j\in J}$ be frame sequences for $\mathcal{H}$ and $\mathcal{K}$, respectively. The synthesis operators T_Ψ and T_Φ map sequences of frame coefficients back to functions in the respective Hilbert space. With this, out

of a discrete realization $u : \ell^2(I) \to \ell^2(J)$ of the operator U one can construct an operator mapping from $\mathcal{H}$ to $\mathcal{K}$ through:

$$T_\Phi \circ u \circ T_\Psi^\dagger : \mathcal{H} \to \mathcal{K}.$$

Operator aliasing error measures the discrepancy between U and its discrete implementation u:

$$\epsilon(U, u, \Psi, \Phi) = U - T_\Phi \circ u \circ T_\Psi^\dagger,$$

with the corresponding scalar error given by the operator norm $\|\epsilon(U, u, \Psi, \Phi)\|$. If this error is zero, the operator U can be perfectly represented by the discrete map u, ensuring continuous-discrete equivalence (CDE).

Representation equivalent Neural Operator (ReNO). To introduce our framework of *Representation equivalent Neural Operators (ReNOs)*, we need to make one more extension. On the level of discretizations, we will consider maps u that take in any pair of frame sequences (Ψ, Φ) of $\mathcal{H}, \mathcal{K}$ and output a mapping from $\ell^2(I)$ to $\ell^2(J)$:

$$u(\Psi, \Phi) : \ell^2(I) \to \ell^2(J).$$

This emphasizes that in order to eliminate aliasing errors, the discrete maps *must depend* on the choice of the discrete representations Ψ and Φ. We then say that a pair (U, u) is a ReNO if, for any pair (Ψ, Φ) with

$$\operatorname{Dom} U \subseteq \mathcal{M}_\Psi \text{ and } \operatorname{Ran} U \subseteq \mathcal{M}_\Phi \tag{1}$$

there is no aliasing error, i.e.

$$\epsilon(U, u, \Psi, \Phi) = 0.$$

Here, $\mathcal{M}_\Psi$ and $\mathcal{M}_\Phi$ are the closed spans of the frame sequences Ψ and Φ. The conditions in (1) simply ensure that only those frame sequences are considered, that can actually represent the domain and range of U, respectively. This definition now straightforwardly extends to a layerwise instantiation: for a neural operator U of the form

$$U = U_L \circ U_{L-1} \circ \cdots \circ U_1,$$

where each U_ℓ is a mapping between separable Hilbert spaces $\mathcal{H}_\ell$ and $\mathcal{H}_{\ell+1}$, each layer must satisfy the ReNO condition (i.e. zero aliasing error). This ensures that any two discrete representations of the operator are equivalent, preserving the underlying structure in function spaces. This framework also allows to introduce ϵ-ReNOs, where a small, controlled amount of aliasing is permissible:

$$\|\epsilon(U, u, \Psi, \Phi)\| \leq \epsilon,$$

for all Ψ and Φ satisfying (1).

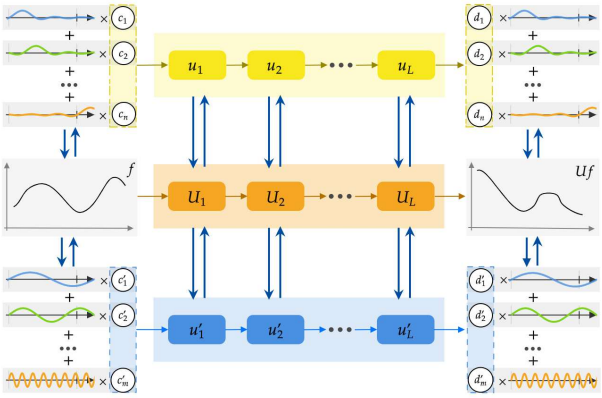

FIGURE 1. A sketch of the ReNO framework. Any learned discretization is tied to the operators in function space through analysis and synthesis operators of frame sequences.

Examples. We define the Representation Equivalence Error as a measure to quantify how severely an architecture deviates from the ReNO property. For $u = u(\Psi, \Phi)$ and $u' = u(\Psi', \Phi')$ one can compute the error between these two discretizations via

$$\tau(u, u') = u - T_{\Phi}^{\dagger} \circ T_{\Phi'} \circ u' \circ T_{\Psi'}^{\dagger} \circ T_{\Psi}.$$

Classical Convolutional Neural Networks (CNNs) and Fourier Neural Operators (FNOs) do not satisfy the ReNO conditions due to inconsistencies across different discretizations. The discrete convolution operation in CNNs does not preserve the continuous-discrete equivalence, leading to aliasing errors. For FNO, the issue lies in the nonlinear activation function: the considered function space is that of bandlimited functions throughout the network. When a pointwise nonlinearity σ, such as the ReLU or GeLU is applied to a function f, then in general, $\sigma(f)$ is no longer bandlimited and has an effective bandwidth significantly larger than that of f.

In contrast, Convolutional Neural Operators (CNOs) take this issue into account: before applying a nonlinear activation function, the function is upsampled (to twice the bandwidth). The application of σ is followed by a downsampling operation to control the dimensionality throughout the network. Convolutional neural operators take the form of a U-Net architecture and convolutions are implemented in the original space without passing to the Fourier domain (see [3] for details and tests on a wide range of benchmark PDEs).

Empirically analyzing the Representation Equivalence Error in different architectures demonstrates that CNNs and FNOs exhibit significant aliasing errors when changing the resolution of input and output frames, while CNOs maintain representation equivalence within a defined resolution range. This highlights the

practical importance of the ReNO framework in ensuring consistency and stability in operator learning.

Discussion. The ReNO framework provides a robust solution to the inconsistencies in neural operator learning by ensuring that discrete representations are equivalent to their continuous counterparts. This framework not only addresses existing challenges but also offers a foundation for developing new neural operators. Future work may explore quantitative measures of error in ReNOs and extend the concept to ϵ-ReNOs, providing a nuanced approach to operator learning with controlled aliasing.

A question that remains open and highly relevant is that of identifying combinations of suitable function spaces and activation functions so that the space is closed under the action of the activation function. More generally, identifying pairs of function spaces B_1, B_2 and activation functions σ for which $\sigma(B_1) \subseteq B_2$ could be helpful in designing new architectures that satisfy the ReNO property.

References

[1] F. Bartolucci, E. de Bézenac, B. Raonić, R. Molinaro, S. Mishra, & R. Alaifari, *Representation equivalent neural operators: a framework for alias-free operator learning*, Advances in Neural Information Processing Systems, 37 (2023).

[2] N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, & A. Anandkumar, *Neural operator: Learning maps between function spaces with applications to PDEs*, Journal of Machine Learning Research, 24(89) (2023), 1-97.

[3] B. Raonić, R. Molinaro, T. De Ryck, T. Rohner, F. Bartolucci, R. Alaifari, S. Mishra, & E. de Bézenac, *Convolutional neural operators for robust and accurate learning of PDEs*, Advances in Neural Information Processing Systems, 37 (2023).

Algorithmic approaches to recovering sparse and low-rank matrices

Johannes Maly

(joint work with Massimo Fornasier, Christian Kuemmerle, Valeriya Naumova)

We consider the reconstruction of sparse[1], low-rank matrices $\mathbf{X}_\star \in \mathbb{R}^{n_1 \times n_2}$ from incomplete and inaccurate measurements

$$\mathbf{y} = \mathcal{A}(\mathbf{X}_\star) + \boldsymbol{\eta} \in \mathbb{R}^m,$$

where $\mathcal{A}\colon \mathbb{R}^{n_1 \times n_2} \to \mathbb{R}^m$ resembles a linear measurement process and $\boldsymbol{\eta} \in \mathbb{R}^m$ models additive noise. This problem, which stems from compressed sensing [2] and related fields, is relevant in several modern applications such as sparse phase retrieval, blind deconvolution of sparse signals, machine learning, and data mining [4, 3, 7]. We examined two algorithmic approaches to the problem:

[1] A vector $\mathbf{z} \in \mathbb{R}^n$ is called s-sparse if at most s of its entries are non-zero. On matrices sparsity can be defined and counted in various ways, cf. [6].

1. Iteratively Re-weighted Least-Squares (IRLS)

The classical iterative IRLS-algorithm for sparse recovery aims at solving an ℓ_1-minimization over the set of data interpolations by iteratively solving weighted ℓ_2-minimization problems via

$$\begin{aligned}\mathbf{x}_{k+1} &= \arg\min_{\mathbf{z}\in\mathbb{R}^n} \langle \mathbf{z}, \mathrm{diag}(\mathbf{w}_k)\mathbf{z}\rangle, \text{ s.t. } \mathbf{A}\mathbf{z} = \mathbf{y} \\ \varepsilon_{k+1} &= \min\left\{\varepsilon_k, \frac{\text{Best s-term approximation error of } \mathbf{x}_{k+1}}{n}\right\} \\ (w_{k+1})_i &= \frac{1}{\max\{|(x_{k+1})_i|, \varepsilon_{k+1}\}}.\end{aligned}$$

To recover the sparse and low-rank matrix $\mathbf{X}_\star$ from $\mathbf{y}$, one only has to replace

- the smoothing parameters ε_k by two sequences

  $$\varepsilon_{k+1} = \min\{\varepsilon_k, \sigma_{R+1}(\mathbf{X}^{(k)})\} \quad \text{and} \quad \delta_{k+1} = \min\{\delta_k, \rho_{s+1}(\mathbf{X}^{(k)})\}$$

- the weights $\mathbf{w}_k$ by the weight operator

  $$W_{\mathbf{X}^{(k)},\varepsilon_k,\delta_k}(\mathbf{Z}) = W^{lr}_{\mathbf{X}^{(k)},\varepsilon_k}(\mathbf{Z}) + \mathbf{W}^{sp}_{\mathbf{X}^{(k)},\delta_k}\mathbf{Z},$$

  where

  $$W^{lr}_{\mathbf{X}^{(k)},\varepsilon_k}(\mathbf{Z}) = \left[\mathbf{U}\&\mathbf{U}_\perp\right]\mathbf{\Sigma}^{-1}_{\varepsilon_k}\begin{bmatrix}\mathbf{U}^* \\ \mathbf{U}^*_\perp\end{bmatrix}\mathbf{Z}\left[\mathbf{V}\&\mathbf{V}_\perp\right]\mathbf{\Sigma}^{-1}_{\varepsilon_k}\begin{bmatrix}\mathbf{V}^* \\ \mathbf{V}^*_\perp\end{bmatrix}$$

  with $\mathbf{\Sigma}_{\varepsilon_k} = \mathrm{diag}(\max\{\sigma_i^{(k)}/\varepsilon_k, 1\})$ and $\mathbf{W}^{sp}_{\mathbf{X}^{(k)},\delta_k} \in \mathbb{R}^{n_1\times n_1}$ with

  $$\left(\mathbf{W}^{sp}_{\mathbf{X}^{(k)},\delta_k}\right)_{ii} = \max\left(\left\|(\mathbf{X}^{(k)})_{i,:}\right\|_2^2/\delta_k^2, 1\right)^{-1}, \qquad \text{for all } i \in [n_1].$$

In [5], we show that under standard restricted isometry property (RIP) assumptions on $\mathcal{A}$ the modified IRLS algorithm exhibits a locally quadratic convergence rate to $\mathbf{X}_\star$ if initialized sufficiently close. Notably, the algorithm automatically balances between sparsity and low-rankness of the iterates. Empirical evaluations furthermore suggest that the obtained bounds on the initialization radius are overly pessimistic in the ambient dimension $n_1 n_2$ of the problem. A major open question of this work is whether the analysis could be refined and whether one can show global convergence of the algorithm under suitable assumptions on $\mathcal{A}$.

2. Multi-penalty optimization

Our second study considers the multi-penalty objective $J^R_{\boldsymbol{\alpha},\boldsymbol{\beta}}\colon \mathbb{R}^{n_1\times R}\times\mathbb{R}^{n_2\times R} \to \mathbb{R}$ defined as

$$J^R_{\boldsymbol{\alpha},\boldsymbol{\beta}}(\mathbf{U},\mathbf{V}) \coloneqq \left\|\mathbf{y} - \mathcal{A}\left(\mathbf{U}\mathbf{V}^\top\right)\right\|_2^2 + \mathrm{EN}_{\boldsymbol{\alpha}}(\mathbf{U}) + \mathrm{EN}_{\boldsymbol{\beta}}(\mathbf{V}),$$

where the elastic net is given by

$$\mathrm{EN}_{\boldsymbol{\gamma}}(\mathbf{Z}) = \gamma_1\|\mathbf{Z}\|_F^2 + \gamma_2\|\mathbf{Z}\|_1,$$

for $\gamma_1, \gamma_2 > 0$. The idea is to regularize the rank by representing the argument as a product of lower dimensional matrices, and to regularize sparsity by ℓ_1-type

regularizers on the single matrix factors. The functional $J^R_{\boldsymbol{\alpha},\boldsymbol{\beta}}$ can be viewed as a natural generalization of the Sparse PCA formulation in [7]. In [1, 6], we analyze the general properties of global minimizers of $J^R_{\boldsymbol{\alpha},\boldsymbol{\beta}}$, propose a non-orthogonal signal model for sparse and low-rank matrices that is tailored to the specific shape of $J^R_{\boldsymbol{\alpha},\boldsymbol{\beta}}$, and introduce robust injectivity to understand the sample complexity that guarantees good approximation of $\mathbf{X}_\star$ by global minimizers of $J^R_{\boldsymbol{\alpha},\boldsymbol{\beta}}$. We prove local convergence of alternating schemes to global minimizers and validate the performance of such approaches in empirical studies. A major open question in this setting is again the exact convergence radius of the numerical optimization methods.

References

[1] M. Fornasier, J. Maly, V. Naumova, *Robust Recovery of Low-Rank Matrices with Non-Orthogonal Sparse Decomposition from Incomplete Measurements*, Applied Mathematics and Computation **392** (2021): 125702.

[2] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, Applied and Numerical Harmonic Analysis (New York: Birkhäuser, 2013).

[3] M. Iwen, A. Viswanathan, Y. Wang, *Robust Sparse Phase Retrieval Made Easy*, Applied and Computational Harmonic Analysis **42**, no. 1 (2017): 135–42.

[4] M. V. Klibanov, P. E. Sacks, and A. V. Tikhonravov, *The Phase Retrieval Problem*, Inverse Problems **11**, no. 1 (1995): 1.

[5] C. Kümmerle and J. Maly, *Recovering Simultaneously Structured Data via Non-Convex Iteratively Reweighted Least Squares*, Advances in Neural Information Processing Systems **36** (2023): 71799–833.

[6] J. Maly, *Robust Sensing of Low-Rank Matrices with Non-Orthogonal Sparse Decomposition*, Applied and Computational Harmonic Analysis **67** (2023): 101569.

[7] H. Zou, T. Hastie, R. Tibshirani, *Sparse Principal Component Analysis*, Journal of Computational and Graphical Statistics **15**, no. 2 (2006): 265–86.

Machine learned regularisation for inverse problems — the dos and don'ts

Carola Schönlieb

Machine learning, in particular deep learning, has entered the field of inverse problems as a prominent and promising new technique, in particular for the ability of deep neural networks to characterise information intrinsic in data to high accuracy. this ability is interesting in the context of regularization. The promises and pitifalls of machine-learned regularization were the topic of this talk, demonstrated on the example of sparse and limited-angle computed tomography.

References

[1] S. Lunz, O. Öktem, C. B. Schönlieb, *Adversarial regularizers in inverse problems*, Advances in neural information processing systems, 31 (2018).

[2] S. Arridge, P. Maass, O. Öktem, C. B. Schönlieb, *Solving inverse problems using data-driven models*, Acta Numerica, 28 (2019), 1-174.

[3] S. Mukherjee, M. Carioni, O. Öktem, C. B. Schönlieb, *End-to-end reconstruction meets data-driven regularization for inverse problems*, Advances in Neural Information Processing Systems, 34 (2021), 21413-21425.
[4] S. Mukherjee, A. Hauptmann, O. Öktem, M. Pereyra, C. B. Schönlieb, *Learned reconstruction methods with convergence guarantees: A survey of concepts and applications*, IEEE Signal Processing Magazine, 40(1) (2023), 164-182.
[5] A. Hauptmann, S. Mukherjee, C. B. Schönlieb, F. Sherry, *Convergent regularization in inverse problems and linear plug-and-play denoisers*, to appear in FOCM (2024).
[6] Z. Shumaylov, J. Budd, S. Mukherjee, C. B. Schönlieb, *Weakly Convex Regularisers for Inverse Problems: Convergence of Critical Points and Primal-Dual Optimisation*, to appear in ICML (2024).

Regularization properties of noise injection

Anna Shalova
(joint work with André Schlichting and Mark A. Peletier)

In an overparametrized setting (when number of parameters of the model is larger than the number of data point) machine learning models have many global minimizers. Not all of the minimizers are equally *good* as some may lead to better generalization. Whether the obtained solution of the optimization problem will generalize well or not may depend on many factors including the choice of the optimization algorithm. In particular, an algorithm can give preference to some minimizers but not the others, this effect is often called *implicit bias*.

In our work [1] we study the implicit bias of *noisy gradient descent* systems: given a gradient descent for some loss function $L : \mathbb{R}^k \to [0, \infty)$

$$w_{k+1} = w_k - \alpha \nabla_w L(w_k),$$

we consider it's noisy version of the form

$$w_{k+1} = w_k - \alpha \nabla_w \hat{L}(w_k, \eta_k), \tag{1a}$$

$$\eta_{k,i} \sim \rho(\sigma), \tag{1b}$$

such that $\hat{L} : \mathbb{R}^{k+d} \to [0, \infty)$ satisfies $\hat{L}(w, 0) = L(w)$ and ρ is a probability distribution with zero mean and variance σ^2. To mention some examples, our theory applies to gradient descent with dropout noise, label noise and minibatching. The noisy gradient descent can have a different behavior compared to the noiseless version, see Figure 1. In particular, under certain assumptions the noisy gradient descent system quickly converges to the set of global minimizers $\Gamma = \{w \in \mathbb{R}^k : L(w) = 0\}$ and then slowly evolves along this set. This slow drift is a result of the stochastic nature of the algorithm and was rigorously characterized for the gradient systems with additive noise in the work of Li et al. [2]. We extend the result of Li et al. to a more general class of system and show that the structure of the noise injection affects not only the drift itself but the time-scale at which it appears.

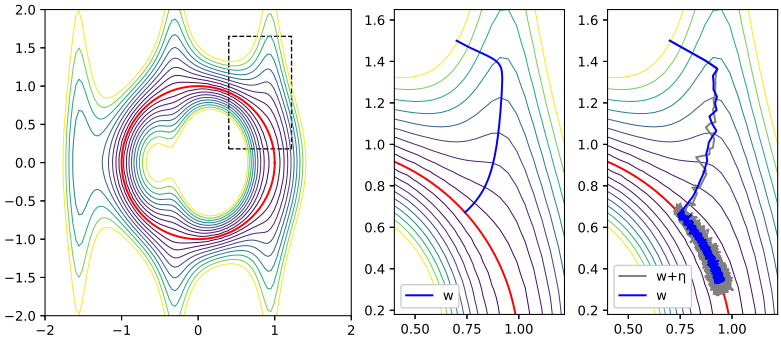

FIGURE 1. Numerical example: the left-hand panel shows the level curves of the loss function L, with the zero-level set Γ marked in red. The middle panel shows a gradient-descent evolution, starting at the top, and converging to Γ. The right-hand panel shows an evolution of the noisy gradient descent (1) with $\hat{L}(w,\eta) := L(w+\eta)$.

To get an intuition of the limiting process assume that dynamics of the parameters w is slow compared to the updates of the noise variables η such that we can average over η:

$$w_{t+1} = w_t - \alpha \mathbb{E}_\eta \big[\nabla_w \hat{L}(w_t, \eta_t)\big].$$

If the noise is small $\sigma^2 \ll 1$, using Taylor expansion we get

$$\begin{aligned} w_{t+1} &= w_t - \alpha \mathbb{E}_\eta \nabla_w \hat{L}(w_k, \eta_k) \\ &= w_t - \alpha \Big(\mathbb{E}_\eta \big[\nabla_w \hat{L}(w_k, 0)\big] + \mathbb{E}_\eta \big[\nabla^2_{w_k,\eta} \hat{L}(w_k, 0)\eta_k\big] \\ &\qquad + \frac{1}{2} \mathbb{E}_\eta \big[\nabla_w \nabla^2_\eta \hat{L}(w_k, 0)[\eta_k, \eta_k]\big] + o(\sigma^2) \Big). \end{aligned}$$

Because for every k the noisy variables η_k are independent random variables with zero mean we have $\mathbb{E}\eta_k = 0$ and $\mathbb{E}_\eta [\eta_k \otimes \eta_k] = \sigma^2 I$, thus giving the resulting dynamic of the form

$$w_{k+1} - w_k = -\alpha \nabla_w L(w_k) - \frac{\alpha}{2} \sigma^2 \nabla_w \Delta_\eta \hat{L}(w_k, 0) + o(\sigma^2) \qquad \text{as } \sigma^2 \to 0\,.$$

We consider a sequence of noisy gradient descent systems which differ by step size α_n and the noise variance σ_n. We make the above observation rigorous by showing that in the limit $\alpha_n, \sigma_n \to 0$ after accelerating the evolution by a factor $1/\alpha_n \sigma_n^2$ the sequence converges to a solution of a gradient flow of the driving functional $\frac{1}{2}\Delta_\eta \hat{L}(w, 0)$. Here we state an informal version of the main result.

Theorem 1. *Assume $\alpha_n \to 0$ and $\sigma_n \to 0$. Set*

$$W_n(\alpha_n \sigma_n^2 k) := w_k \,.$$

Then the sequence W_n converges to a limit curve $W = (W(t))_{t>0}$, that satisfies the constrained gradient flow

$$\partial_t W(t) = -\frac{1}{2} P_\Gamma \nabla_w \Delta_\eta \hat{L}(W(t), 0),$$

where P_Γ is the orthogonal projection onto the tangent plane of Γ at $W(t)$.

However for some noisy loss functions $\hat{L}$ that are often used on practice, the dynamic given by Theorem 1 is trivial while experiments indicate otherwise. This is, for example, the case for L_2 function approximation loss function with label noise. To explain this possible mismatch we extend our analysis to the class of noisy loss function having quadratic behaviour in η and yielding a trivial result according to Theorem 1. Namely we study $\hat{L}$ of form

$$\hat{L}(w, \eta) = L(w) + f(w) \cdot \eta + \tfrac{1}{2} H(w) : (\eta \otimes \eta) + g(\eta), \tag{2}$$

for certain smooth maps f, H, and g. We assume that $g(0) = 0$ and that each diagonal element H_{ii} vanishes, so that $\Delta_\eta \hat{L}(w, 0) = \sum_{i=1}^{d} H_{ii}(w) = 0$.

It turns out that in such a case nontrivial dynamic appears at a slower time scale $1/\alpha^2\sigma^2$. The limiting process is still restricted to the set of global minimizers but is now a stochastic process. Moreover, for convergence it is enough to consider the case of inifinitely small step-size $\alpha_n \to 0$ with any finite noise variance $\sigma_n \to \sigma_0$. So the following theorem holds:

Theorem 2. *Let $\hat{L}$ be a degenerate loss function as described above. For $\alpha_n \to 0$ and $\sigma_n \to \sigma_0 \geq 0$, let $(w_k^n)_{k\geq}$ be the noisy gradient descent and set*

$$W_n(\alpha_n^2 \sigma_n^2 k) := w_k \,.$$

Then the sequence W_n converges to a limit $W = (W(t))_{t>0}$ that satisfies the constrained stochastic differential equation

$$dW(t) = P_\Gamma \beta(W(t)) dB_t + F(W(t)) : \beta(W(t))\beta(W(t))^\top dt,$$

with $W(t) \in \Gamma$ for $t > 0$. Here β and F are given in terms of f and H, and B is a multidimensional standard Brownian motion.

Applying our results to a few examples we show that:

(1) *Dropout* noise induces non-trivial dynamics at the fast time scale $1/\alpha\sigma^2$,
(2) *Label* noise leads to a deterministic evolution but at a slower time-scale $1/\alpha^2\sigma^2$,
(3) *Minibatch* noise has a trivial dynamics at both of the discussed time-scales if not combined with additional ways of noise injection.

In addition, both dropout and label noise limiting processes can be interpreted as a type of weight-regularization.

References

[1] A. Shalova, A. Schlichting and M. Peletier, *Singular-limit analysis of gradient descent with noise injection*, arXiv:2404.12293 (2024).

[2] Z. Li, T. Wang and S. Arora, *What Happens after SGD Reaches Zero Loss?–A Mathematical Framework*, International Conference on Learning Representations (2022).

Breaking the quadratic bottleneck in non-convex matrix sensing: Near-optimal recovery guarantees with linear rank-dependency

Dominik Stöger

(joint work with Yizhe Zhu)

Low-rank matrix recovery problems are ubiquitous in many areas of science and engineering. Most of the methods that have been studied for these problems can roughly be divided into two categories: Convex optimization approaches based on nuclear norm minimization, and non-convex approaches that use factorized gradient descent.

While methods from the latter category are typically computationally much less expensive, basically all existing recovery guarantees for factorized gradient descent are more pessimistic with respect to the number of samples required. In particular, they require the number of samples to scale quadratically with the rank of the ground truth matrix. This raises the question whether one can obtain recovery guarantees for the non-convex methods if the samples size scales linear in the degrees of freedom.

In this talk we consider the scenario that one obtains m observations of the form

$$y_i = \langle \mathbf{A}_i, \mathbf{X}_\star \rangle := \text{trace } (\mathbf{A}_i \mathbf{X}_\star) \qquad \text{for } i = 1, 2, \dots, m.$$

Here, $(\mathbf{A}_i)_{i=1}^n \subset \mathbb{R}^{d\times d}$ represent known and symmetric measurement matrices whose entries are i.i.d. with standard Gaussian distribution $\mathcal{N}(0,1)$ on the diagonal and $\mathcal{N}(0,1/2)$ on the off-diagonal. The goal is to recover the unknown ground truth matrix $\mathbf{X}_\star \in \mathbb{R}^{d\times d}$ which we assume to be of rank r and, moreover, *symmetric and positive definite*.

For that, we follow a two-stage approach, which was originally proposed by Keshavan, Montanari, and Oh [2] and which is widely studied in the literature. It consists of the following two steps.

(1) **Spectral initialization:** Let $\mathbf{M} = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top$ be the truncated rank-r singular value decomposition of the matrix $\frac{1}{m}\sum_{i=1}^m \langle \mathbf{A}_i, \mathbf{X}_\star\rangle \mathbf{A}_i$. Then the initialization is defined as $\mathbf{U}_0 := \mathbf{V}\mathbf{\Sigma}^{1/2} \in \mathbb{R}^{d\times r}$.

(2) **Gradient descent:** We consider the non-convex objective

$$f(\mathbf{U}) := \frac{1}{m}\sum_{i=1}^m \left(y_i - \langle \mathbf{A}_i, \mathbf{U}\mathbf{U}^\top\rangle\right)^2,$$

where $\mathbf{U} \in \mathbb{R}^{d\times r}$. The objective is minimized via

$$\mathbf{U}_t = \mathbf{U}_{t-1} - \mu \nabla f\left(\mathbf{U}_{t-1}\right) \qquad \text{for } t = 1, 2, \ldots$$

Here, $\mu > 0$ represents the step size.

To state our main results, we need to introduce a few definitions. The condition number of the rank-r matrix $\mathbf{X}_\star$ is defined via

$$\kappa := \frac{\lambda_1\left(\mathbf{X}_\star\right)}{\lambda_r\left(\mathbf{X}_\star\right)}.$$

Moreover, let $\mathbf{M}_\star \in \mathbb{R}^{d\times r}$ such that $\mathbf{X}_\star = \mathbf{M}_\star \mathbf{M}_\star^\top$. Then, we can define the following notion of distance

$$\operatorname{dist}\left(\mathbf{U}_t, \mathbf{M}_\star\right) := \min_{\mathbf{R}\ \text{rotation}} \|\mathbf{U}_t\mathbf{R} - \mathbf{M}_\star\|_F$$

Now we are prepared to state our main result.

Theorem 1. *Assume that the sample size satisfies $m \gtrsim rd\kappa^4$. Moreover, assume that the step size satisfies $\mu \asymp \frac{1}{\kappa\|\mathbf{X}_\star\|}$. Then, with high probability, it holds that*

$$dist\left(\mathbf{U}_t, \mathbf{M}_\star\right) \lesssim r\left(1 - c\mu\sigma_{\min}(\mathbf{X}_\star)\right)\big)^t \sqrt{\lambda_r(\mathbf{X}_\star)},$$

where $c > 0$ is some absolute constant.

This result shows that non-convex methods based on factorized gradient descent can recover the ground truth matrix as soon as the number of measurements scales linearly with the degrees of freedom of the ground truth matrix.

The talk is based on the article [1], which is currently in preparation.

References

[1] D. Stöger, Y. Zhu, *Breaking the quadratic rank-bottleneck in non-convex matrix sensing: Recovery guarantees with optimal sample complexity*, in preparation.

[2] R. Keshavan, A. Montanari, and S. Oh, *Matrix completion from a few entries*, IEEE Trans. Inf. Theory 56, No. 6 (2010), 2980–2998.

Stability and Sampling results for gabor phase retrieval

Philipp Grohs

We consider the problem of reconstructing $f \in L^2$ from (samples of) the spectrogram

$$|V_\varphi f(w,\omega)|^2 = \left| \int f(t)\varphi(t-x)\exp(-2\pi i\omega t)dt \right|^2.$$

Our main results are twofold:

(1) The stability of this reconstruction is governed by the degree of disconnectedness of $|V_\varphi f|^2$. This is quantified in terms of a Cheeger/Poincaré type constant.

(2) Sampling on a regular lattice cannot yield unique recovery but one can get uniqueness using judiciously chosen irregular sampling points.

Mathemalchemy: a mathematical and artistic adventure

Ingrid Daubechies

Mathemalchemy is a collaborative art installation conceived as the brainchild of mathematician and physicist Ingrid Daubechies and fiber artist Dominique Ehrmann, and driven by the energy and enthusiasm of 24 mathematical artists and artistic mathematicians. The installation celebrates the creativity and beauty of mathematics. Playful constructs include a flurry of Koch snowflakes, Riemann basalt cliffs, and Lebesgue terraces. It was designed and constructed during the pandemic, and has been touring North America since January 2022; it will soon move to its 5th exhibition venue. The talk will review the genesis and creation of the installation, and highlight some of its mathematical features.

The following videos were shown during the talk:

- "Bakery – Through the Mathemalchemy Looking Glass"
- "Happy Pi Day!"

Bilipschitz invariants

Dustin G. Mixon

(joint work with Jameson Cahill, Joseph W. Iverson, Daniel Packer, Yousef Qaddura)

Most data processing algorithms are designed for Euclidean data, and so data scientists are inclined to represent objects in Euclidean space. Unfortunately, such a representation can introduce ambiguity. For example, a graph on n vertices might be represented by its adjacency matrix, which resides in the space of real symmetric $n \times n$ matrices, but the same graph can be represented by any member of its orbit under the conjugation action of S_n. As such, the naïve Euclidean distance between representations fails to capture a notion of distance between the underlying objects.

More generally, we consider the setting in which our data resides in a Hilbert space V, and we identify points modulo a group G of linear isometries. In this setting, it's appropriate to consider the *quotient distance* defined by

$$d(x, y) := \inf_{\substack{p \in G \cdot x \\ q \in G \cdot y}} \|p - q\|_V.$$

This determines a quotient metric space $V/\!/G$, whose points are the topological closures of the G-orbits in V. In order to make use of the data science tools in Euclidean space, we seek a bilipschitz embedding of the quotient $V/\!/G$ into some Hilbert space. Furthermore, in order to preserve the signal from our metric, we wish to minimize *distortion*, that is, the quotient of the optimal upper and lower Lipschitz bounds of our bilipschitz map.

Let's start by considering some examples. If $V = \mathbb{R}^d$ and G is a reflection group, then we can achieve distortion 1 by mapping each orbit to its intersection with a

fixed Weyl chamber. If $V = \mathbb{R}^d$ and $G \leq O(d)$ is finite, then for generic templates $z_1, \ldots, z_n \in \mathbb{R}^d$, the map

$$x \mapsto \left\{ \max_{g \in G} \langle g z_i, x \rangle \right\}_{i=1}^n$$

descends to a bilipschitz map $\mathbb{R}^d /\!/ G \to \mathbb{R}^n$ provided $n \geq 2d$; this can be seen by combining the injectivity result from [3] with the fact that injectivity implies bilipschitz for this map [1]. (In some cases, like if G is a reflection group, one may take n to be even smaller [4, 5].) Furthermore, one may estimate the distortion of this map for random templates. The remainder of our discussion is based on [2].

The case of $V = \mathbb{R}$ and $G = \{\pm 1\}$ is a simple instance of both examples considered above, in which case one is inclinded to consider the bilipschitz map determined by $x \mapsto |x|$. By contrast, algebraic invariant theory leads one to consider the G-invariant polynomials, namely the even polynomials, which are generated as an algebra by the polynomial x^2. Hilbert proved that more generally, for sufficiently nice groups $G \leq GL(d)$ (e.g., finite groups), the G-invariant polynomials separate all G-orbits and are finitely generated as an algebra, and so taking any finite generator set as coordinate functions produces an injective G-invariant map. However, in the case of $V = \mathbb{R}$ and $G = \{\pm 1\}$, the map determined by $x \mapsto x^2$ is not lower Lipschitz when x is small. In fact, whenever G is finite, this difficulty emerges for any differentiable invariant near points with a nontrivial stabilizer. In cases where the origin is the only point with a nontrivial stabilizer, one may interpret generating polynomials as coordinate functions on the sphere and then homogeneously extend to produce a bilipschitz embedding for the whole space. For example, if $V = \mathbb{R}^d$ and $G = \{\pm \mathrm{id}\}$, then the invariant polynomials are generated by $\{x_i x_j\}_{i \leq j}$, and the map

$$x \mapsto \frac{x \otimes x}{\|x\|}$$

is simply $x \mapsto |x|$ when $d = 1$, and is bilipschitz with distortion $\sqrt{2}$ when $d \geq 2$.

The *Euclidean distortion* $c_2(X)$ of a metric space X is the minimum possible distortion of a bilipschitz map from X into an arbitrary Hilbert space. It turns out that $c_2(X)$ equals the supremum of $c_2(F)$ over finite sub-metric spaces F of X, and furthermore, computing any such $c_2(F)$ reduces to a semidefinite program. With this reduction, one may show that

$$c_2\big(\mathbb{R}^d /\!/ \{\pm \mathrm{id}\}\big) = \sqrt{2}$$

when $d \geq 2$, meaning the embedding described above exhibits the smallest possible distortion. We are also interested in the Euclidean distortion of other quotient metric spaces, especially

$$\ell^2(\mathbb{N}; \mathbb{R}^2) /\!/ S_\infty \qquad \text{and} \qquad \ell^2(\mathbb{Z}; \mathbb{R}) /\!/ \mathbb{Z}$$

due to their relationship with other quotient metric spaces. These are currently open questions.

References

[1] R. Balan, E. Tsoukanis, *G-invariant representations using coorbits: Bi-Lipschitz properties*, 2308.11784 (2023).
[2] J. Cahill, J. W. Iverson, D. G. Mixon, *Towards a bilipschitz invariant theory*, Appl. Comput. Harmon. Anal., to appear.
[3] J. Cahill, J. W. Iverson, D. G. Mixon, D. Packer, *Group-invariant max filtering*, Found. Comput. Math. (2004).
[4] D. G. Mixon, D. Packer, *Max filtering with reflection groups*, Adv. Comput. Math. **49** (2023).
[5] D. G. Mixon, Y. Qaddura, *Injectivity, stability, and positive definiteness of max filtering*, arXiv:2212.11156 (2022).

Covariance estimation under one-bit quantization

Sjoerd Dirksen
(joint work with Johannes Maly, Holger Rauhut)

A common task in signal processing is to estimate the correlation matrix or the covariance matrix of a high-dimensional Gaussian distribution from i.i.d. samples that have been quantized to finitely many bits. In my talk I will consider a setup where each entry of each sample is quantized to one bit using an efficient, memoryless quantizer. In this setup, a well-known approach in the engineering literature is to use the arcsin law (also known as Grothendieck's identity) to estimate the correlation matrix. I will present non-asymptotic, near-optimal error bounds for this type of estimator in terms of the spectral norm. Surprisingly, the bounds reveal that this estimator can outperform the sample covariance matrix (of the samples before quantization) in certain scenarios. I will also show that by using dithering, i.e., adding well-designed noise before quantization, one can estimate the full covariance matrix of any subgaussian distribution from quantized samples at the same (minimax optimal) rate as the sample covariance matrix. This second result is based on a new version of the Burkholder-Rosenthal inequalities for matrix martingales. The talk is based on [1, 2].

References

[1] S. Dirksen, J. Maly, H. Rauhut, *Covariance estimation under one-bit quantization*, Annals of Statistics **50.6** (2022), 3538–3562.
[2] S. Dirksen, J. Maly, *Tuning-free one-bit covariance estimation using data-driven dithering*, IEEE Transactions on Information Theory, to appear, arXiv:2307.12613

A primer on physics-informed machine learning

Claire Boyer
(joint work with Francis Bach, Gérard Biau, Nathan Doumèche)

Physics-informed machine learning combines the expressiveness of data-based approaches with the interpretability of physical models. In this context, we consider a general regression problem where the empirical risk is regularized by a partial differential equation that quantifies the physical inconsistency. More formally, given a sample $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$ of i.i.d. copies of (X, Y), the goal is to construct an estimator $\hat{f}_n$ of $f^\star$ based on these n observations. The distinctive element of PIML is the inclusion of a prior on $f^\star$, asserting its compliance with a known PDE. Therefore, it is assumed that $f^\star$ is at least weakly differentiable, belonging to the Sobolev space $H^s(\Omega)$ for some integer $s > d/2$, and that there is a known differential operator $\mathcal{D}$ such that $\mathcal{D}(f^\star) \simeq 0$. For instance, if the desired solution $f^\star$ is intended to conform to the wave equation, then $\mathcal{D}(f)(x,t) = \partial^2_{t,t} f(x,t) - \partial^2_{x,x} f(x,t)$ for $(x,t) \in \Omega$. Overall, we are interested in the minimizer of the empirical risk function

$$R_n(f) = \frac{1}{n}\sum_{i=1}^{n} |f(X_i) - Y_i|^2 + \lambda_n \|f\|^2_{H^s(\Omega)} + \mu_n \|\mathcal{D}(f)\|^2_{L^2(\Omega)} \tag{1}$$

We prove that for linear differential priors, the problem can be formulated as a kernel regression task, giving a rigorous framework to analyze physics-informed ML. In particular, the physical prior can help in boosting the estimator convergence (compared to the classical Sobolev kernel estimator).

References

[1] N. Doumèche, F. Bach, C. Boyer, and G. Biau, *Physics-informed machine learning as a kernel method*, arXiv (2024).

Curvature on graphs

Stefan Steinerberger

This talk is concerned with various notions of curvature on combinatorial graphs $G = (V, E)$, in particular:

(1) combinatorial notions derived from the Gauss–Bonnet theorem
(2) analytic notions derived from Optimal Transport, in particular, the Olivier–Ricci curvature and the Lin–Lu–Yau curvature
(3) and a potential-theoretic notion that is defined as follows: if $D \in \mathbb{R}^{n\times n}$ is the graph distance matrix, meaning

$$D_{ij} = d(v_i, v_j),$$

then the curvature in vertex v_i is define as the $i-$th entry of $D^{-1}\mathbf{n}$, where

$$\mathbf{n} = (n, n, \dots, n) \quad \text{is the constant vector.}$$

Note that $D^{-1}\mathbf{n}$ may not be uniquely defined or not defined at all. If it is not defined at all, then one may take the pseudo-inverse to recover much of the subsequent theory. However, for reasons that are still unexplained, it tends to be defined (though not necessarily uniquely). However, one inherits a type of uniqueness.

Proposition 1. *Let G be a connected graph and suppose $Dw_1 = \boldsymbol{n} = Dw_2$ for two vectors $w_1, w_2 \in \mathbb{R}^n_{\geq 0}$. Then $\|w_1\|_{\ell^1} = \|w_2\|_{\ell^1}$.*

One can now establish a result in the style of Bonnet–Myers: lower bounds on the curvature imply upper bounds on the diameter.

Theorem 1. *Let G be connected and suppose $Dw = \boldsymbol{n}$. If $w_i \geq K$,*

$$diam(G) \leq \frac{2n}{\|w\|_{\ell^1}} \leq \frac{2}{K}.$$

This mirrors Bonnet–Myers-type results for other notions of curvature. This result, in turn, follows from a much more general result that can be derived from the von Neumann Minimax Theorem. More details can be found in [2]. The entire approach appears to be robust in the sense that D can be replaced by other notions of distance, see [1].

References

[1] K. Devriendt, A. Ottolini, and S. Steinerberger, *Graph curvature via resistance distance*, Discrete Applied Mathematics **348**, (2024), 68–78.

[2] S. Steinerberger, *Curvature on graphs via equilibrium measures*, Journal of Graph Theory **103**, (2023), 415–436.

Generalized Norm Resolvent Convergence and Stability for Graph Convolutional Networks

Christian Koke

Graph Convolutional Networks (GCNs) are a prominent class of machine learning architectures adapted to operating on graph structured data. Using the standard graph Laplacian $L \in \mathbb{R}^{N\times N}$ and node-feature matrix $X \in \mathbb{R}^{N\times F}$ to represent an N-node graph G together with its F-dimensional node feature vectors $\{X_{i:}\}_{1\leq i\leq N}$, a (graph level) graph convolutional network Φ generates a Euclidean embedding $\vec{F} \in \mathbb{R}^d$ in a d-dimensional representation space for each such graph G as

$$\vec{F} = \Phi(L, X). \tag{1}$$

This is done by first iteratively updating the node feature matrix X and finally running a (graph-isomorphism invariant) aggregation scheme combining node-wise information into a graph-level representation. The individual steps $\mathbb{R}^{N\times F_{\ell-1}} \ni X^{\ell-1} \mapsto X^{\ell} \in \mathbb{R}^{N\times F_\ell}$ in the iterative update of X are implemented (ignoring bias terms for simplicity) as

$$X^{\ell}_{:i} = \mathrm{ReLU}\left(\sum_{j=1}^{F_{\ell-1}} h^{\ell}_{\theta_{ij}}(L) X^{\ell-1}_{:j}\right). \tag{2}$$

Here the set of matrices $\{h^{\ell}_{\theta_{ij}}(L)\}_{ij}$ arises from the application of learned functions $\{h^{\ell}_{\theta_{ij}}(\cdot)\}_{ij}$ to the Laplacian L. The ultimate graph level representation $\vec{F}$ is then obtained by calculating the $\|\cdot\|_1$-norm for each column of the final representation $X^{\ell_{\text{final}}} \in \mathbb{R}^{N\times d}$.

Given two distinct graphs $G, \tilde{G}$ that are defined on different node sets and which possess distinct adjacency structures but nevertheless describe the same underlying object, we want to ensure the feature vectors $\vec{F}, \vec{\tilde{F}}$ generated for these two graphs by the same graph neural network Φ are similar in the sense that

$$\left\|\vec{F} - \vec{\tilde{F}}\right\| \ll 1. \tag{3}$$

Such a setting might e.g. occur if the two graphs arise from one-another via coarse graining, rewiring or can be taken to discretize the same underlying manifold.

In this talk we discuss how filters $\{h^{\ell}_{\theta_{ij}}(\cdot)\}_{ij}$ in (2) need to be chosen so that (3) can be be guaranteed. To this end, we make use of the concept of generalized norm resolvent convergence [1] and consider two graphs $G, \tilde{G}$ to be close if the resolvents $R_{-1}(L) := (L + Id)^{-1}$ of their respective Laplacians satisfy

$$\|R_{-1}(L) - \tilde{J}R_{-1}(\tilde{L})J\| \ll 1. \tag{4}$$

Here J and $\tilde{J}$ are linear intertwining operators, with J mapping from $\ell^2(G)$ to $\ell^2(\tilde{G})$ and $\tilde{J}$ mapping in the opposite direction.

We then establish that we can bound the difference of the generated graph embeddings $\vec{F}, \vec{\tilde{F}}$ as

$$\left\|\vec{F} - \vec{\tilde{F}}\right\| \equiv \left\|\Phi(X, L) - \Phi(JX, \tilde{L})\right\| \lesssim \|X\| \cdot \|R_{-1}(L) - \tilde{J}R_{-1}(\tilde{L})J\| \tag{5}$$

precisely if learned filter functions $\{h^{\ell}_{\theta_{ij}}(\cdot)\}_{ij}$ arise as Laplace transforms of finite complex measures on the positive real line. More precisely, we need to demand that each such h may be represented as

$$h(z) = \int_0^{\infty} e^{-tz} d\mu(t) \tag{6}$$

for some finite complex Borel measure on the positive real line $\mathbb{R}_{\geq 0} = [0, \infty)$ for which $\mu(\{0\}) = 0$.

Using this result, we then discuss how stable graph neural networks that are able to deal with noisy or multi-scale data can be designed. Finally we investigate the performance capabilities and stability properties of such networks in real-world experiments.

References

[1] O. Post, *Spectral Analysis on Graph-like Spaces*, Lecture Notes in Mathematics **2039** (2012), Springer.

On the frame set property of Hermite functions

Markus Faulhuber
(joint work with Irina Shafkulovska, Ilya Zlotnikov)

We denote the the Hilbert space of square-integrable functions on the line by $L^2(\mathbb{R})$ and its inner product and norm by $\langle .\ ,\ .\rangle$ and $||.||$, respectively.

We are interested in the frame set property of Gabor systems with a Hermite function. A Gabor system with *window* $g \in L^2(\mathbb{R})$ (non-zero) over a *lattice* $\Lambda \subset \mathbb{R}^2$ is a structured function system of the form

$$\mathcal{G}(g,\Lambda) = \{M_\omega T_x g \mid (x,\omega) \in \Lambda \subset \mathbb{R}^2\}.$$

Here, T_x and M_ω denote the unitary operators of translation (time-shift) and modulation (frequency-shift), respectively:

$$T_x g(t) = g(t-x) \quad \text{and} \quad M_\omega g(t) = e^{2\pi i\omega t}, \quad t,x,\omega \in \mathbb{R}.$$

The composition is called a time-frequency shift, denoted by

$$\pi(z) = M_\omega T_x, \quad z = (x,\omega) \in \mathbb{R}^2.$$

Note that T_x and M_ω do not commute in general, which also implies that time-frequency shifts do not commute in general. We have the commutation relation $M_\omega T_x = e^{2\pi i\omega x} T_x M_\omega$.

A Gabor system is a frame for $L^2(\mathbb{R})$, if and only if there exist positive constants $0 < A \le B < \infty$, called *frame bounds*, such that

$$A||f||^2 \le \sum_{\lambda\in\Lambda} |\langle f, \pi(\lambda)g\rangle|^2 \le B||f||^2, \quad \forall f \in L^2(\mathbb{R}).$$

In this case any function has a stable expansion with respect to $\mathcal{G}(g,\Lambda)$ of the form

$$f = \sum_{\lambda\in\Lambda} c_\lambda \pi(\lambda) g, \quad (c_\lambda) \in \ell^2(\Lambda).$$

The n-th order Hermite function h_n is the Gaussian function multiplied with the n-th order Hermite polynomial H_n, subject to normalization:

$$h_n(t) = H_n(t)e^{-\pi t^2} = c_n(-1)^n e^{\pi t^2}\frac{d^n}{dt^n}e^{-2\pi t^2},\ t\in\mathbb{R}, \quad \text{so} \quad ||h_n|| = 1,\ n\in\mathbb{N}_0,$$

with $c_n = 2^{1/4}/\sqrt{n!(2\pi)^n 2^n}$. Since the window is a fixed Hermite function, we seek to find lattice parameters $(a,b) \in \mathbb{R}_+ \times \mathbb{R}_+$ for the rectangular lattice

$$\Lambda_{(a,b)} = a\mathbb{Z} \times b\mathbb{Z}$$

such that the resulting Gabor system $\mathcal{G}(h_n, \Lambda_{(a,b)})$ is a frame.

Known obstructions and results. The Balian-Low theorem [1, 6] in combination with the density condition for Gabor systems gives the following result.

If $ab \ge 1$, then the Gabor system $\mathcal{G}(h_n, \Lambda_{(a,b)})$ cannot be a frame.

In the other direction, we have a sufficiency result of Gröchenig and Lyubarskii [4].

If $ab < \frac{1}{n+1}$, then the Gabor system $\mathcal{G}(h_n, \Lambda_{(a,b)})$ is a frame.

For growing n, the gap between the known obstructions and the sufficiency result clearly becomes large. Further known exceptions, due to Lyubarskii and Nes [7] as well as counter-examplex provided by Lemvig [5], lead to the following result for the square lattice $\Lambda_{\left(\frac{1}{\sqrt{n+1}},\frac{1}{\sqrt{n+1}}\right)} = \frac{1}{\sqrt{n+1}}\mathbb{Z} \times \frac{1}{\sqrt{n+1}}\mathbb{Z} = \frac{1}{\sqrt{n+1}}\mathbb{Z}^2$.

The Gabor system $\mathcal{G}(h_n, \frac{1}{\sqrt{n+1}}\mathbb{Z}^2)$ is not a frame for $n = 0, 1, 2, 3$.

One is tempted to see a pattern which may continue, but once $n \geq 4$ our main result from [2] shows that the opposite holds true[1].

The Gabor system $\mathcal{G}(h_n, \frac{1}{\sqrt{n+1}}\mathbb{Z}^2)$ is a frame for $n \geq 4$.

Proof idea. The result is based on the *Janssen representation* of the *Gabor frame operator*. The frame operator is acting on functions $f \in L^2(\mathbb{R})$ by the rule

$$S_{g,\Lambda} f = \sum_{\lambda \in \Lambda} \langle f, \pi(\lambda) g \rangle \, \pi(\lambda) g.$$

Assume g is "nice" (e.g., coming from Feichtinger's algebra). By considering the inner products $\langle S_{g,\Lambda} f, h \rangle$ and an application of the Poisson summation formula in combination with the orthogonality relations (a version of the Moyal identity), one arrives at

$$S_{g,\Lambda} = \mathrm{vol}(\mathbb{R}^2/\Lambda)^{-1} \sum_{\lambda^\circ \in \Lambda^\circ} \langle g, \pi(\lambda^\circ) g \rangle \; \pi(\lambda^\circ),$$

which is the Janssen representation of $S_{g,\Lambda}$. Here Λ° is the symplectic dual or adjoint lattice, which is characterized by

$$\Lambda^\circ = \{\lambda^\circ \in \mathbb{R}^2 \mid \pi(\lambda)\pi(\lambda^\circ) = \pi(\lambda^\circ)\pi(\lambda), \; \forall \lambda \in \Lambda\}.$$

In fact, it is simply a 90 degrees rotated version of the classical dual lattice. Note:

$$\Lambda^\circ_{(a,b)} = \frac{1}{b}\mathbb{Z} \times \frac{1}{a}\mathbb{Z}, \quad \text{in particular} \quad \left(\frac{1}{\sqrt{n+1}}\,\mathbb{Z}^2\right)^\circ = \sqrt{n+1}\,\mathbb{Z}^2$$

With the Janssen representation at hand, we can use the *Janssen test* [8], [9].

Let g be such that $||g|| = 1$. If $\sum_{\lambda^\circ \in \Lambda^\circ} |\langle g, \pi(\lambda^\circ) g \rangle| < 2$, then $\mathcal{G}(g, \Lambda)$ is a frame.

The condition follows by comparing $\mathrm{vol}(\mathbb{R}^2/\Lambda) S_{g,\Lambda}$ to the identity operator on $L^2(\mathbb{R})$ and using a Neumann argument for the invertibility of the frame operator.

Our proof then makes use of the *Laguerre connection* [3] between Hermite functions and Laguerre functions. Denote the n-th order Laguerre polynomial by

$$\mathcal{L}_n(t) = \sum_{k=0}^{n} \binom{n}{k} \frac{(-t)^k}{k!}.$$

Then, we have the following formula

$$|\langle h_n, \pi(z) h_n \rangle| = \left|\mathcal{L}_n\left(\pi \; |z|^2\right)\right| \; e^{-\frac{\pi}{2}|z|^2}, \quad z = (x, \omega) \in \mathbb{R}^2.$$

[1] — WHAT?! —

By rigorous analysis and new and established bounds on Laguerre polynomials, we show in [2] that

$$\sum_{(k,l)\in\mathbb{Z}^2} \left|\mathcal{L}_n(\pi(n+1)(k^2+l^2))\right| e^{-\frac{\pi}{2}(n+1)(k^2+l^2)} < 2, \quad n \geq 4.$$

This is, by the Laguerre connection and the Janssen test, sufficient for the Gabor system $\mathcal{G}(h_n, \frac{1}{\sqrt{n+1}}\mathbb{Z}^2)$ to be a frame once $n \geq 4$.

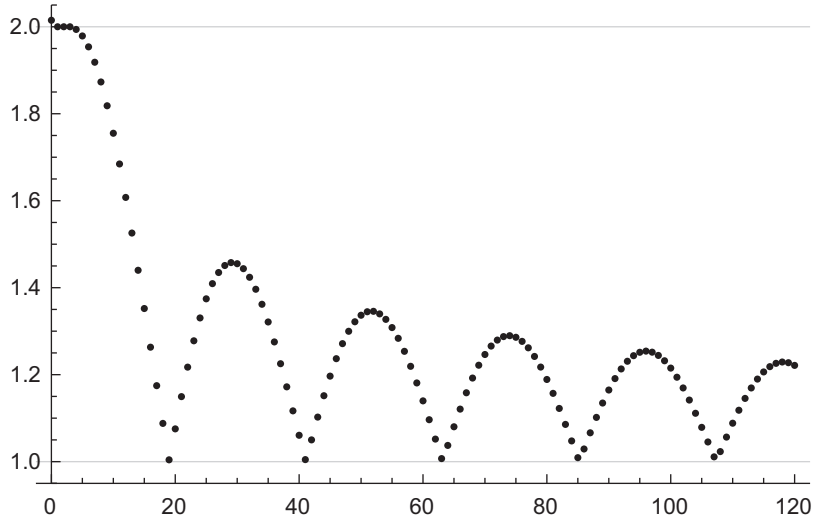

FIGURE 1. Some values of the Janssen test for $\mathcal{G}(h_n, \frac{1}{\sqrt{(n+1)}}\ \mathbb{Z}^2)$.

REFERENCES

[1] R. Balian, *Un principe d'incertitude fort en théorie du signal ou en mécanique quantique*, Comptes rendus de l'Académie des sciences, Série II, 292(20) (1981), 1357–1361.

[2] M. Faulhuber, I. Shafkulovska, and I. Zlotnikov, *On the frame property of Hermite functions and analytic extensions of their frame sets*, arXiv preprint (2024), 2403.10503.

[3] G. B. Folland, *Harmonic analysis in phase space*, Number 122 in Annals of Mathematics Studies, Princeton University Press (1989).

[4] K. Gröchenig and Y. Lyubarskii, *Gabor frames with Hermite functions*, Comptes Rendus Mathematique (2007), 344(3):157–162.

[5] J. Lemvig, *On some Hermite series identities and their applications to Gabor analysis*, Monatshefte für Mathematik (2017), 182(4):889–912.

[6] F. E. Low, *Complete sets of wave packets*, A passion for physics – Essays in honor of Geoofrey Chew, World Scientific, Singapore (1985), pp. 17–22.

[7] Y. Lyubarskii and P. G. Nes, *Gabor frames with rational density*, Applied and Computational Harmonic Analysis (2013), 34(3):488–494.

[8] T. Tschurtschenthaler, *The Gabor Frame Operator (its Structure and Numerical Consequences)*, Master's thesis, University of Vienna (2000).

[9] C. Wiesmeyr, *Construction of frames by discretization of phase space*, PhD thesis, University of Vienna (2013).

Can quantum harmonic analysis explain structure in data? The example of data augmentation

Monika Dörfler
(joint work with Franz Luef, Henry McNulty)

Data, which we strive to use, interpret, classify, depend on some kind of representation. That means, that each data point is written in some dictionary of building blocks by means of coefficients. For high-dimensional, complex data sets, structured dimensionality reduction methods are essential in order to enable useful further processing [1]. The guiding idea behind some approaches is the hypothesis, that learning is made possible by the fact that data of interest live on low-dimensional manifolds as opposed to the dimensionality of the space in which the data are collected a priori. Inspired by this idea, our approach hinges on the idea that the dimension of data is not at all canonical, but can vary according to the chosen representation of the data. Ideally, the choice of representation avoids loss of essential information, where the latter may depend on the application at hand. When dealing with time-series of any kind, time-frequency methods are often applied to obtain image-like representations of time-series data such as speech and music and to encode the impact of variance over time. Applying convolutional neural network (CNN) architectures to the resulting TF-transformed versions of time-series data points has been surprisingly successful in various machine learning (ML) tasks. The underlying processes are, however, not entirely understood. One important hypothesis is the assumption that the informative content of the data actually lies on a manifold of significantly lower dimensions than their domain. It is not clear, however, how these essential parts of data can be made explicit. In our work, we propose to use quantum harmonic analysis tools for the identification of time-frequency localized components that determine the entropy of a data set. In particular, certain intrinsic structures, which repeat over a data set of interest, may be encoded as TF-local components. Intuitively, these components determine the coefficients of the convolutional kernels in the lower levels of the network and thus can be expected to carry the essential structure for a certain problem at hand. Thus, the TF-local components are considered useful for the identification of an underlying data manifold.

We make a connection between data augmentation of original time-domain data and their *effective dimensionality* [2] via tools from quantum harmonic analysis developed in [3, 4, 5]. The key insight guiding these efforts is the association of the data operator $S_{\mathcal{D}}$ to a system of functions $\mathcal{D}$. This allows us to capture the structure of the data and their interaction. Data augmentation is then formalized as the mixed-state localization operator corresponding to $S_{\mathcal{D}}$. Since augmentation can be interpreted as a generalized convolution in the quantum harmonic analysis context, we use existing results [6, 7] that show that operators corresponding to augmented data sets yield smoother principle components than the original data sets. This provides a mathematical explanation of the benefits of augmentation,

since the principal components determine an approximation of the data set and increased smoothness is desirable and suppresses over-fitting effect.

Future work and open questions include the hypothesis, that correct augmentation augments along data manifold. Can we possibly learn an underlying data manifold from data by using concepts of local entropy and local correlation? Can this help to improve augmentation strategies? Is it possible to use the concept of non-linear principal components in order tofully exploit the manifold hypothesis?

Detailed information on the technical background may be found in [9, 8].

References

[1] S. T. Roweis, L. K. Saul, *Nonlinear Dimensionality Reduction by Locally Linear Embedding*, Science (2000), 290(5500):2323–2326.

[2] O. Roy, M. Vetterli, *The effective rank: A measure of effective dimensionality*, In 15th European Signal Processing Conference (2007), 606–610.

[3] F. Luef, E. Skrettingland, *Convolutions for localization operators*, J. Math. Pures Appl. (9) (2018), 118:288–316.

[4] F. Luef, E. Skrettingland, *Mixed-state localization operators: Cohens class and trace class operators*, Journal of Fourier Analysis and Applications (2019), 25(4): 2064–2108.

[5] F. Luef, E. Skrettingland, *On accumulated Cohens class distributions and mixed-state localization operators*, Constr Approx (2020), 52:3164.

[6] Á. Bényi, K.A. Okoudjou, *Modulation Spaces*, In: Modulation Spaces. Applied and Numerical Harmonic Analysis (2020).

[7] K. Gröchenig, C. Heil, *Modulation Spaces as Symbol Classes for Pseudodifferential Operators* (2003).

[8] M. DoerflerF. Luef,E. Skrettingland, *Local Structure and effective Dimensionality of Time Series Data Sets*, arXiv 2111.02153 (2023).

[9] M. Dörfler, F. Luef, H. McNulty, E. Skrettingland, *Time-frequency analysis and coorbit spaces of operators*, J. Math. Anal. Appl. 534, 2 (2024), 128058 MR:MR4685895.

Optimal sampling for stochastic gradient descent

Philipp Trunschke

(joint work with Robert Gruhlke, Charles Miranda, Anthony Nouy)

Consider the problem of minimising a *cost functional*

$$\operatorname*{minimise}_{v\in\mathcal{M}} \mathcal{L}(v), \qquad \mathcal{L}(v) := \int \ell(v;x)\,\mathrm{d}\rho(x)$$

over a nonlinear *model class* of functions $\mathcal{M} \subseteq L^2(\rho)$. When ρ is a probability measure, and only realisations of the loss are accessible, the exact objective must be replaced with a Monte Carlo estimate before standard first-onder methods like gradient descent can be employed. This results in the well-known *stochastic gradient descent* (SGD) method. However, replacing the true objective with an estimate ensues a "generalisation error". Rigorous bounds for this error usually require L^∞-compactness of $\mathcal{M}$ and Lipschitz continuity of ℓ while providing only a very slow decay with increasing sample size. This slow decay is unfavourable in settings where high accuracy is required, or sample creation is costly.

	GD	Best-case	Worst-case	SGD
L-smoothness	$\mathcal{O}(t^{-1})$	$\mathcal{O}(t^{-1+\varepsilon})$	$\mathcal{O}(t^{-1/2+\varepsilon})$	$\mathcal{O}(t^{-1/2+\varepsilon})$
strong convexity	$\mathcal{O}(a^t)$	$\mathcal{O}(a^t)$	$\mathcal{O}(t^{1-2\varepsilon})$	$\mathcal{O}(t^{1-2\varepsilon})$

TABLE 2. Almost sure convergence rates for different algorithms with $\varepsilon \in (0, \frac{1}{2})$ and $a \in (0, 1)$ depending on the chosen step size.

To address this issue, we draw inspiration from the linear least squares problem, where optimal sampling methods can be used to obtain quasi-optimal minimisers. Although the quasi-optimality results do not extend to general, nonlinear model classes, we can leverage their existence in the linear case by performing a sequence of updates in local linearisations of the model class.

To be specific, we suppose that in every step $t \in \mathbb{N}$ there exists a linear space $\mathcal{T}_t$ that approximates $\mathcal{M}$ locally around the current iterate u_t. Given the gradient $g_t := \nabla\mathcal{L}(u_t) \in L^2(\rho)$ and a quasi-optimal estimate P_t^n of the $L^2(\rho)$-orthogonal projector P_t onto $\mathcal{T}_t$, we perform a linear update $\bar{u}_{t+1} := u_t - s_t P_t^n g_t$ in direction of the (empirically) projected negative gradient $-P_t^n g_t$. This yields the intermediate iterate $\bar{u}_{t+1} \in L^2(\rho)$. Since the $\bar{u}_{t+1}$ is not guaranteed to lie in the original model class $\mathcal{M}$, we perform a recompression step $u_{t+1} := R_t(\bar{u}_{t+1})$, where $R_t : L^2(\rho) \to \mathcal{M}$ takes the linear update $\bar{u}_{t+1}$ back to the model class $\mathcal{M}$ with a controllable error in the cost $\mathcal{L}$. The proposed algorithm can thus be presented in the two equations

$$\bar{u}_{t+1} := u_t - s_t P_t^n g_t, \qquad g_t := \nabla\mathcal{L}(u_t),$$
$$u_{t+1} := R_t(\bar{u}_{t+1}).$$

Under classical assumptions on the cost functional $\mathcal{L}$ and the sequences of projectors P_t^n, step sizes s_t and recompressions R_t, the proposed optimisation scheme converges almost surely to a stationary point of the true objective. The corresponding convergence rates are displayed in Table 2. We find that the proposed algorithm exhibits the same convergence rates as classical *gradient descent* (GD) in the best case but can never perform worse than SGD. We pay particular attention to the estimation of the projectors P_t^n, which must be carried out using optimally weighted samples to achieve the presented rates.

Implicit bias of policy gradient in linear quadratic control: extrapolation to unseen initial states

NOAM RAZIN

(joint work with Yotam Alexander, Edo Cohen-Karlik, Raja Giryes, Amir Globerson, Nadav Cohen)

In modern machine learning, models can often fit training data in numerous ways, some of which perform well on unseen (test) data, while others do not. Remarkably, in such cases gradient descent frequently exhibits an implicit bias that leads to excellent performance on unseen data. This implicit bias was extensively studied

in supervised learning, but is far less understood in optimal control, which in a broad sense is equivalent to reinforcement learning. There, learning a controller applied to a system via gradient descent is known as policy gradient, and a question of prime importance is the extent to which a learned controller extrapolates to unseen initial states. In this talk, I will present a recent work [1] that theoretically studies the implicit bias of policy gradient in terms of extrapolation to unseen initial states. Focusing on the fundamental Linear Quadratic Regulator (LQR) problem, we establish that the extent of extrapolation depends on the degree of exploration induced by the system when commencing from initial states included in training. Experiments corroborate our theory, and demonstrate its conclusions on problems beyond LQR, where systems are non-linear and controllers are neural networks. We hypothesize that real-world optimal control may be greatly improved by developing methods for informed selection of initial states to train on.

References

[1] N. Razin, Y. Alexander, E. Cohen-Karlik, R. Giryes, A. Globerson, N. Cohen, *Implicit Bias of Policy Gradient in Linear Quadratic Control: Extrapolation to Unseen Initial States*, arXiv preprint 2402.07875

L^1 matrix norms, gauges and factorizations

Radu Balan
(joint work with Fushuai Jiang)

In this talk we consider the decomposition of positive semidefinite matrices as a sum of rank one matrices. We introduce and investigate the properties of various measures of optimality of such decompositions. For some classes of positive semidefinite matrices we give explicitly these optimal decompositions. Motivated by a question raised by H. Feichtinger at an Oberwolfach workshop in 2004 (exactly 20 years ago!) we formulate the following problem. For each integer n, find the optimal constant $C(n) \geq 1$ such that for any positive semidefinite matrix $A \in \mathbb{C}^{n\times n}$,

$$\gamma_+(A) := \inf_{A=\sum_{k=1}^m x_k x_k^*} \sum_{k=1}^m ||x_k||_1^2 \leq C(n) \sum_{i,j=1}^n |A_{i,j}|.$$

It turns out that one can choose $m = n^2$ so the infimum is achieved. We show the following properties:

(1) The map $A \mapsto \gamma_+(A)$ is continuous over the closed cone of psd matrices.
(2) The optimum value is given by the following infinite dimensional linear program:

$$\gamma_+(A) = \min_{\substack{\mu \in B(S_1) \\ A = \int_{S_1} xx^Y d\mu(x)}} \mu(S_1),$$

where $B(S_!)$ denotes the convex set of Borel measures over the unit sphere S_1 with respect to l^1 norm in $\mathbb{C}^n$.

(3) The same constant $C(n)$ satisfies, for every $T = T^* \in \mathbb{C}^{n\times n}$:

$$\max_{||x||_1\leq 1} < Tx, x > \leq C(n) \max_{\substack{A = A^* \geq 0 \\ \sum_{i,j} |A_{ij}| \leq 1}} trace(TA).$$

(4) The dual of the last SDP problem has strong duality gap:

$$\max_{\substack{A = A^* \geq 0 \\ \sum_{i,j} |A_{ij}| \leq 1}} trace(TA) = \min_{S=S^*\geq 0} \max_{i,j} |T_{ij} + S_{ij}|.$$

References

[1] R. Balan, K. Okoudjou, A. Poria, *On a Feichtinger Problem*, Operators and Matrices vol. **12**(3), 881-891 (2018).

[2] R. Balan, K.A. Okoudjou, M. Rawson, Y. Wang, R. Zhang, *Optimal l1 Rank One Matrix Decomposition*, Harmonic Analysis and Applications, Rassias M., Ed. Springer (2021).

Sampling numbers of the Fourier-analytic Barron spaces

Felix Voigtlaender

Originally introduced and studied in the 90s [1, 2, 3], the space of *Barron functions* has recently received a lot of attention [5, 7, 8, 11] in the wake of the deep learning revolution [12], due to its close connection to neural networks.

We consider Barron functions on the d-dimensional unit cube $\Omega_d := [-\frac{1}{2}, \frac{1}{2}]^d$. Precisely, a function $f : \Omega_d \to \mathbb{R}$ is called a *(Fourier-analytic) Barron function with smoothness* $\alpha > 0$ if there exists a (measurable) function $F : \mathbb{R}^d \to \mathbb{C}$ with

$$(1) \qquad \begin{aligned} f(x) &= \int_{\mathbb{R}^d} F(\xi) \cdot e^{2\pi i \langle x,\xi\rangle}\, d\xi \quad \text{for all } x \in \Omega_d, \\ \text{and} \qquad \|F\|_{L^1_\alpha} &:= \int_{\mathbb{R}^d} (1+|\xi|)^\alpha \cdot |F(\xi)|\, d\xi < \infty. \end{aligned}$$

The unit ball U^α_d of the Barron space consists of all such functions f for which one can choose F such that $\|F\|_{L^1_\alpha} \leq 1$.

In the papers [1, 2, 3], Barron showed for the case $\alpha = 1$ that functions in U^1_d can be well approximated by *shallow neural networks*, without suffering from the *curse of dimensionality*. In modern terminology, let $\varrho : \mathbb{R} \to \mathbb{R}$, $\varrho(x) = \max\{0, x\}$ denote the *ReLU activation function* [10]. Then a shallow neural network (or more precisely, a *shallow ReLU network*) with N neurons is a function of the form

$$\Phi_N : \quad \mathbb{R}^d \to \mathbb{R}, \qquad \Phi_N(x) = \sum_{j=1}^N c_j\, \varrho(\langle w_j, x\rangle + b_j)$$

for certain parameters $c_j, b_j \in \mathbb{R}$ and $w_j \in \mathbb{R}^d$, which are collectively called the *weights of the network*. The main approximation result derived by Barron [1] (in a slightly modified form, taken from [5]) reads as follows:

Theorem 1. *There exists a universal constant $\kappa > 0$, such that for each $f \in U_d^1$ and each $N \in \mathbb{N}$, there exists a shallow ReLU network Φ_N with N neurons and all weights bounded in absolute value by κ such that $\|f - \Phi_N\|_{L^\infty} \leq \kappa \cdot \sqrt{d/N}$.*

As noted in [3], this implies that given m random samples $(X_i, f(X_i))$ of an unknown Barron function $f \in U_d^1$ (with X_i uniformly distributed in $\Omega_d = [-\frac{1}{2}, \frac{1}{2}]^d$), one can reconstruct f up to L^2 error of expected size $\mathcal{O}\big((\frac{d\log(m)}{m})^{1/4}\big)$.

In this talk, we present results from the preprint [14] concerning the *optimal rate* of reconstructing an unknown function $f \in U_d^\alpha$ from m point samples $(x_i, f(x_i))_{i=1}^m$ as $m \to \infty$. Here, the location of the sampling points x_i can be chosen freely (also using randomness, *if desired*) to allow for optimal reconstruction. In the language of *information-based complexity*, this means that we aim to determine the asymptotic behavior of the so-called *(non-linear) sampling numbers* of the space of Barron functions with smoothness $\alpha > 0$, formally given by

$$\sigma_m(U_d^\alpha)_{L^p} := \inf_{x_1,\dots,x_m \in \Omega_d} \inf_{R:\mathbb{R}^m \to L^p(\Omega_d)} \sup_{f \in U_d^\alpha} \big\| f - R\big(f(x_1), \dots, f(x_m)\big) \big\|_{L^p}.$$

This question was originally raised in [4, Page 34]. Our main result is as follows:

Theorem 2 ([14])**.** *There exist constants $C_1, C_2 > 0$ depending only on d and α such that for all $m \in \mathbb{N}$ and $p \in [1, \infty]$, we have*

$$C_1 \cdot m^{-(\frac{1}{\max\{p,2\}} + \frac{\alpha}{d})} \leq \sigma_m(U_d^\alpha)_{L^p} \leq C_2 \cdot [\ln(e+m)]^{2+8\frac{\alpha}{dp}} \cdot m^{-(\frac{1}{\max\{p,2\}} + \frac{\alpha}{d})}.$$

Remark (Open questions).

① In the literature, certain other types of spaces[1] are also called *Barron spaces* (sometimes *variational Barron spaces*); see e.g. [7]. The (metric) entropy numbers and the nonlinear dictionary approximation rates of these spaces have been studied in [13], but the question of the asymptotic behavior of the (non-linear) sampling numbers for these spaces is still open, to the best of the knowledge of the present author.

② Ignoring logarithmic factors, the above estimate shows that the sampling numbers exhibit the asymptotic behavior $\sigma_m(U_d^\alpha)_{L^p} \asymp m^{-(\frac{1}{\max\{p,2\}} + \frac{\alpha}{d})}$ as $m \to \infty$ for *fixed* d. The constants C_1, C_2, however, depend in an unspecified way on the input dimension d. It is an open problem to determine the *pre-asymptotic behavior* of the sampling numbers of the Barron spaces, i.e., the behavior of $\sigma_m(U_d^\alpha)_{L^p}$ as a function of d and m.

③ The above result we allows a free choice of the sampling points. From a statistical learning perspective, it is interesting to study the optimal reconstruction error when using *(uniform) random sampling points*. The work [3] shows that one can achieve the error $\mathcal{O}\big((\frac{d\log(m)}{m})^{1/4}\big)$ in this setting, but it is unclear if this is optimal. To the knowledge of the present author, this question is still open.

[1] See [5] for a discussion of inclusion relations between different types of Barron spaces.

Proof ideas. We discuss the proof ideas for the upper bound for the case $p = 2$; for the full argument (also for general $p \in [1, \infty]$) and for the proof of the lower bound (also for randomized algorithms), see [14]. The proof proceeds in three steps:

① Let $\varphi \in C_c^\infty\big((-\frac{1}{2}, \frac{1}{2})^d\big)$ be arbitrary. Then, using standard arguments from Fourier analysis, one can show that for each $f \in U_d^\alpha$, the localized function $\varphi \cdot f$ has a Fourier series with quickly decaying coefficients. More precisely, we have[2]

$$(\varphi \cdot f)(x) = \sum_{n \in \mathbb{Z}^d} c_n e^{2\pi i \langle n, x\rangle} \quad \text{for all } x \in \Omega_d, \qquad \text{with} \qquad \sum_{n \in \mathbb{Z}^d} (1 + |n|)^\alpha |c_n| \lesssim 1.$$

From this, simply by truncating the Fourier series to the N largest coefficients, one sees by a variant of *Stechkin's inequality* (see e.g. [9, Proposition 2.3]) that for each $N \in \mathbb{N}$, there exists a trigonometric polynomial p_N with only N non-zero coefficients satisfying $\|\varphi \cdot f - p_N\|_{L^2} \lesssim N^{-(\frac{1}{2} + \frac{\alpha}{d})}$ and also $\deg p_N \lesssim N^{\frac{1}{d} + \frac{1}{2\alpha}}$.

Remarkably, as shown in [6], a similar estimate also holds for the L^∞ norm instead of the L^2 norm. There thus exists a trigonometric polynomial q_N with $\deg q_N \lesssim N^{\frac{1}{d} + \frac{1}{2\alpha}}$, with only N non-zero coefficients, and $\|\varphi \cdot f - q_N\|_{L^\infty} \lesssim N^{-(\frac{1}{2} + \frac{\alpha}{d})}$.

② Given samples of f, one can generate corresponding samples of $\varphi \cdot f$. *These samples can then be interpreted as noisy samples of the* sparse *trigonometric polynomial q_N from Step* ①. Then, using results from compressive sensing (see [9, Corollary 12.34(b)]), one can obtain a good reconstruction of φf. By choosing $\varphi \equiv 1$ on the smaller cube $\Omega_d^* := [-\frac{1}{4}, \frac{1}{4}]^d$, one can thus obtain a good reconstruction of f on the smaller cube Ω_d^*, if one has access to suitable samples of f on the larger cube $\Omega_d = [-\frac{1}{2}, \frac{1}{2}]^d$. This can be regarded as a *local reconstruction result.*

③ Using a "Whitney-type" decomposition of the cube Ω_d into infinitely many cubes with the size of each cube being proportional to its distance to the boundary of the "big" cube, one can "upgrade" this local reconstruction result to a global one. This uses the fact that, since the error is measured in L^2, in order to achieve a given accuracy, only finitely many of the cubes need to be considered, since the remaining measure will be small. □

References

[1] A. R. Barron, *Neural net approximation*, Proc. 7th Yale workshop on adaptive and learning systems. Vol. 1 (1992).

[2] A. R. Barron, *Universal approximation bounds for superpositions of a sigmoidal function*, IEEE Transactions on Information Theory, vol. 39, no. 3 (1993), pp. 930-945.

[3] A. R. Barron, *Approximation and estimation bounds for artificial neural networks*. Mach. Learn. **14** (1994), 115–133.

[4] P. Binev, A. Bonito, R. DeVore, and G.Petrova, *Optimal learning*, Calcolo **61**, 15 (2024).

[5] A. Caragea, P. Petersen, and F. Voigtlaender, *Neural network approximation and estimation of classifiers with classification boundary in a Barron class*, Ann. Appl. Probab. **33**, no. 4 (2023), 3039–3079.

[6] R. A. DeVore and V. N. Temlyakov, *Nonlinear approximation by trigonometric sums*, J. Fourier Anal. Appl. **2** (1995), no. 1, 29–48.

[2] This is not true in general *without* the localization by φ, since this would entail that f has a *continuous* $\mathbb{Z}^d$-periodic extension to $\mathbb{R}^d$, which is not generally true for Barron functions.

[7] W. E and S. Wojtowytsch, *Representation formulas and pointwise properties for Barron functions*, Calc. Var. Partial Differential Equations **61** (2022), Paper No. 46, 37.
[8] W. E, C. Ma, and L. Wu, *The Barron Space and the Flow-Induced Function Spaces for Neural Network Models*, Constr. Approx. **55** (2022), 369–406.
[9] S. Foucart and H. Rauhut, *A mathematical introduction to compressive sensing*, Applied and Numerical Harmonic Analysis, Birkhäuser/Springer, New York (2013).
[10] K. He, X. Zhang, S. Ren, and J. Sun, *Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification*, Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2015), pp. 1026-1034
[11] J. M. Klusowski and A. R. Barron, *Risk bounds for high-dimensional ridge function combinations including neural networks*, arXiv preprint, arXiv:1607.01434.
[12] Y. LeCun, Y. Bengio, G. Hinton, *Deep learning*, Nature **521** (2015), 436–444.
[13] L. Ma, J. W. Siegel, and J. Xu, *Uniform approximation rates and metric entropy of shallow neural networks*, Res. Math. Sci. **9**, 46 (2022).
[14] F. Voigtlaender, *L^p sampling numbers for the Fourier-analytic Barron space*, arXiv preprint, arXiv:2208.07605.

Johnson-Lindenstrauss Embeddings with Kronecker Structure

Felix Krahmer
(joint work with Stefan Bamberger, Rachel Ward)

The problem of embedding a finite point cloud into a lower-dimensional space in a way that approximately preserves the structure, i.e. their pairwise Euclidean distances, was first studied by Johnson and Lindenstrauss in [1]. Their construction is a random linear map, which is why random matrix contructions with this property are referred to as Johnson-Lindenstrauss (JL) embedding. In distributional form, the JL property reads as follows:

Definition: A random matrix $\Phi \in \mathbb{R}^{m\times N}$ is a Johnson-Lindenstrauss embedding if for any $x \in \mathbb{R}^N$ one has

$$\mathbb{P}(\big|\|\Phi x\|_2^2 - 1\big| > \epsilon) < \eta. \tag{1}$$

JL embeddings have found numerous applications in many fields, such as numerical linear algebra or machine learning. The general idea is typically that if one aims to solve a high-dimensional problem, one can often find an approximate solution by considering the corresponding problem for the low-dimensional data obtained by applying the JL embedding. The combined approach, however, is most advantageous if the embedding can also applied in a fast way. This insight has motivated a line of research on *fast* Johnson-Lindenstrauss embeddings, which aims for constructing and analyzing random matrices Φ with the JL property, which also have structure allowing for fast matrix-vector multliplication. One of the first constructions of this kind was the fast JL transform introduced in [7], in the form of a randomly row-subsampled discrete Hadamard matrix with randomized column signs. This construction was later improved and refined in [2], and ultimately extended to also cover very large point clouds in [3] by establishing a near-equivalence between the Johnson-Lindenstrauss embedding property and a deterministic restricted isometry property [3].

In these considerations, no distinction is made between different classes of signals. Motivated by applications in polynomial sketching [8] and least squares problems with tensor structure [9, 4], recent works have refined this objective aiming for JL embeddings that are partircularly fast on vectors with a tensor structure. For example, [5] proposed the use of a row-subsampled discrete Hadamard matrix with column signs randomized according to a Kronecker-structured Rademacher vector, and conjectured that such an embedding satisfies the Johnson-Lindenstrauss property. More precisely, these constructions are very fast for data points of the form $x = x^{(1)} \otimes \cdots \otimes x^{(d)} \in \mathbb{R}^{n^d} = \mathbb{R}^N$, i.e., Kronecker products of d data vectors, each of dimension n, provided the embedding matrix $\Phi \in \mathbb{R}^{m \times N}$ itself has Kronecker structure $\Phi = \Phi^{(1)} \otimes \cdots \otimes \Phi^{(d)}$, where the dimensions of the factors of Φ correspond to the factor dimensions of x. Indeed, then the matrix-vector multiplication Φx can be factored as $\Phi x = (\Phi^{(1)} x^{(1)}) \otimes \cdots \otimes (\Phi^{(d)} x^{(d)})$, and *can be computed factor by factor, without constructing x explicitly.* Our results reviewed in this presentation (see [6] for the journal article) improve, simplify, and generalize the current embedding results for the Kronecker Johnson-Lindenstrauss embedding [9, 8, 4] by generalizing an approach from [3] on near-equivalence between Johnson-Lindenstrauss property and the restricted isometry property to JL embeddings with Kronecker structure to higher-degree tensor embeddings.

More precisely, we sharpen the existing bounds on the embedding dimension for which a Johnson-Lindenstrauss embedding with Kronecker structure exists to $m = C_d \frac{1}{\epsilon^2} (\log(1/\eta))^d$, up to logarithmic factors in $\log(1/\eta)$, $1/\epsilon$, and in N, improving the results in [8] by a factor of $\log(1/\eta)$. In particular, for the case of $d = 2$ at the core of the oblivious sketching procedure [8], our results improve the scaling of the embedding dimension in $\log(\frac{1}{\eta})$ from cubic to quadratic.

We additionally establish that this embedding result is optimal in the η dependence by providing a lower bound of $m = \Theta((\log(1/\eta))^d)$. We achieve the optimal bounds by generalizing the near-equivalence between the JL property and the restricted isometry property of [3] to higher-order tensors, in a sharper way than what was shown in [9]. The key tool is a concentration inequality for random tensors analogous to the Hanson-Wright inequality, that we developed and proved in [10].

References

[1] W.B. Johnson and J. Lindenstrauss, *Extensions of Lipschitz Mappings into a Hilbert Space*, Contemporary Mathematics **26** (1984), 189–206.

[2] N. Ailon and E. Liberty, *An Almost Optimal Unrestricted Fast Johnson-Lindenstrauss Transform*, ACM Transactions on Algorithms **9**, no. 3 (2013), 1–12

[3] F. Krahmer and R. Ward, *New and Improved Johnson-Lindenstrauss Embeddings via the Restricted Isometry Property*, SIAM Journal on Mathematical Analysis **43**, no. 3 (2011), 1269–1281.

[4] M. Iwen, D. Needell, E. Rebrova, and A. Zare *Lower Memory Oblivious (Tensor) Subspace Embeddings with Fewer Random Bits: Modewise Methods for Least Squares*, SIAM Journal on Matrix Analysis **42**, no. 1 (2021), 376–416.

[5] C. Battaglino, G. Ballard, and T. Kolda, *A practical randomized CP tensor decomposition*, SIAM Journal on Matrix Analysis **39**, no. 2 (2018), 876–901.

[6] S. Bamberger, F. Krahmer, and R. Ward, *Johnson-Lindenstrauss Embeddings with Kronecker Structure*, SIAM Journal on Matrix Analysis **43**, no. 4 (2022), 1806–1850.
[7] N. Ailon and B. Chazelle, *Approximate Nearest Neighbors and the Fast Johnson-Lindenstrauss Transform*, Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing (2006), 557—563.
[8] Thomas D Ahle et al., *Oblivious sketching of high-degree polynomial kernels*, Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (2020), 141–160.
[9] Ruhui Jin, Tamara G. Kolda, and Rachel Ward, *Faster Johnson-Lindenstrauss Transforms via Kronecker Products*, Information and Inference, a Journal of the IMA, **10**, no. 4 (2021), 1533–1562.
[10] S. Bamberger, F. Krahmer, and R. Ward, *The Hanson–Wright inequality for random tensors*, Sampling Theory, Signal Processing and Data Analysis **20**, 14 (2022).

Optimal recovery from inaccurate data

Simon Foucart

Optimal Recovery, a subfield of Approximation Theory, is regaining momentum is the Data Science era, as it proposes a learning theory framework focusing on worst-case scenarios by stepping away from the statistical assumption that the data acquisition process is random. More precisely, an element f from a Banach space F—think of f as a function—needs to be recovered from data available in the form $y = \Lambda f$ for some fixed linear map $\Lambda : F \to \mathbb{R}^m$, called observation map. Additional information translating some prior scientific knowledge is available in the form of a model assumption $f \in \mathcal{K}$, $\mathcal{K}$ being a subset of F. A recovery map for the estimation of a linear map $Q : F \to Z$, landing in a normed space Z, is simply a map $\Delta : \mathbb{R}^m \to Z$. Its performance is assessed through its global worst-case error

$$\sup_{f \in \mathcal{K}} \|Q(f) - \Delta(\Lambda f)\|_Z.$$

The goal of Optimal Recovery is, not surprisingly, to uncover recovery maps with the smallest global worst-case error possible (or close). In realistic situations, however, we need to adjust the ideal setting just described to account for inaccurate data now of the form $y = \Lambda f + e$ for some nonzero $e \in \mathbb{R}^m$. This is discussed in two distinct parts of the presentation.

In a first part, taking place when $F = H$ is a Hilbert space, we model the error vector deterministically by assuming that it belongs to the set $\{e \in \mathbb{R}^m : \|e\|_2 \le \eta\}$, while the set $\mathcal{K}$ is taken as $\{f \in H : \|Pf\| \le \varepsilon\}$ for some linear operator P on H. We reveal that a globally optimal recovery map is constructed with the help of the regularization map

$$\Delta_\tau : y \in \mathbb{R}^m \mapsto \operatorname*{argmin}_{f \in H} \left[(1-\tau)\|Pf\|^2 + \tau\|\Lambda f - y\|_2^2\right] \in H$$

with a parameter τ that can be selected in a principled way. Precisely, an optimal recovery map is given by $Q \circ \Delta_{\tau^\sharp}$ where $\tau^\sharp = d^\sharp/(c^\sharp + d^\sharp)$ and $c^\sharp, d^\sharp$ are solutions to

$$\operatorname*{minimize}_{c,d \ge 0} c\varepsilon^2 + d\eta^2 \qquad \text{subject to} \quad cP^*P + d\Lambda^*\Lambda \succeq 0.$$

This result, as well as extensions covering the cases of mixed accurate/inaccurate data and ℓ_1-inaccurate data, can be deduced (see [2] for full details) from the accurate-data scenario involving the two-hyperellipsoid-intersection model set

$$\mathcal{K} = \{f \in H : \|Rh\| \leq 1 \text{ and } \|Sh\| \leq 1\}.$$

As an open problem, we call for a complete solution in the local optimality setting, too.

In a second part, taking place when $\mathcal{K}$ is a symmetric and convex subset of an arbitrary vector space F but restricting our attention to quantities of interest Q that are linear functionals, we view $e \in \mathbb{R}^m$ as a random vector with log-concave distribution (which includes Gaussian, Laplace, and uniform distributions). The notion of global worst-case error is thus modified to

$$\mathrm{ge}_p^{\mathrm{or}}(\Delta) = \left(\mathbb{E}\left[\sup_{f \in \mathcal{K}} \left|Q(f) - \Delta(\Lambda f + e)\right|^p\right]\right)^{1/p}, \qquad p \in [1, \infty).$$

As shown in [3], linear recovery maps are now near-optimal, in the sense that there exists a constant $\kappa_p > 1$ such that

$$\inf_{\substack{\Delta : \mathbb{R}^m \to \mathbb{R} \\ \Delta \text{ linear}}} \mathrm{ge}_p^{\mathrm{or}}(\Delta) \leq \kappa_p \inf_{\Delta : \mathbb{R}^m \to \mathbb{R}} \mathrm{ge}_p^{\mathrm{or}}(\Delta).$$

This extends a result specific to Gaussian noise obtained by Donoho in [1], who considered instead of $\mathrm{ge}_p^{\mathrm{or}}$ a notion of global worst-case error where the expectation and the supremum were swapped. An open line of inquiries concerns the full recovery problem, i.e., the case $Q = \mathrm{Id_F}$.

References

[1] D. L. Donoho, *Statistical estimation and optimal recovery*, The Annals of Statistics **22** (1994), 238–270.

[2] S. Foucart and C. Liao, *Radius of information for two intersected centered hyperellispoids and implications in Optimal Recovery from inaccurate data*, Journal of Complexity **83** (2024), 101841.

[3] S. Foucart and G. Paouris, *Near-optimal estimation of linear functionals with log-concave observation errors*, Information and Inference **12** (2023), 2546–2561.

Kernel regime of deep neural networks: insights and limitations

Mariia Seleznova

Training dynamics of non-linear Deep Neural Networks (DNNs) are notoriously difficult to study, so the current theory heavily relies on simplifications. Remarkably, DNNs' dynamics simplify dramatically in the infinite-width limit, where DNNs enter the so-called kernel regime under certain conditions. The dynamics in the kernel regime are linearized around the initialization and are governed by deterministic and constant Neural Tangent Kernel (NTK). Thus, optimization and generalization of DNNs in the kernel regime can be studied theoretically using the NTK, and many recent works have adopted this approach. Given that modern DNNs are typically overparametrized, it appears plausible that the infinite-width

limit provides a promising framework for such models. However, many authors have pointed out limitations of this approach. In this talk, we discuss whether the kernel regime provides a good approximation for the behaviour of deep fully-connected networks. Our results reveal that the answer depends on the network's depth-to-width ratio and the distribution of parameters at initialization. While we conclude that deep networks are generally not in the kernel regime at the beginning of training, we also propose a new approach to study DNNs' dynamics using the kernel regime in the end of training.

References

[1] M. Seleznova, G. Kutyniok, *Neural tangent kernel beyond the infinite-width limit: Effects of depth and initialization*, International Conference on Machine Learning (ICML), PMLR (2022).
[2] M. Seleznova, D. Weitzner, R. Giryes, G. Kutyniok, H. H Chou, *Neural (tangent kernel) collapse*, Advances in Neural Information Processing Systems 36 (2024).

Zak-OTFS for integration of sensing and communication

Robert Calderbank
(joint work with Muhammad Ubadah, Saif Khan Mohammed, Ronny Hadani, Shachar Kons, Ananthanarayanan Chockalingam)

The Zak-OTFS input/output (I/O) relation is predictable and non-fading when the delay and Doppler periods are greater than the effective channel delay and Doppler spreads, a condition which we refer to as the crystallization condition. The filter taps can simply be read off from the response to a single Zak-OTFS point (impulse) pulsone waveform, and the I/O relation can be reconstructed for a sampled system that operates under finite duration and bandwidth constraints. Predictability opens up the possibility of a model-free mode of operation. The time-domain realization of a Zak-OTFS point pulsone is a pulse train modulated by a tone, hence the name, pulsone. The Peak-to-Average Power Ratio (PAPR) of a pulsone is about 15 dB, and we describe a general method for constructing a spread pulsone for which the time-domain realization has a PAPR of about 6dB. We construct the spread pulsone by applying a type of discrete spreading filter to a Zak-OTFS point pulsone. The self-ambiguity function of the point pulsone is supported on the period lattice Λ_p, and by applying a discrete chirp filter, we obtain a spread pulsone with a self-ambiguity function that is supported on a rotated lattice Λ^*. We show that if the channel satisfies the crystallization conditions with respect to Λ^* then the effective DD domain filter taps can simply be read off from the cross-ambiguity between the channel response to the spread pulsone and the transmitted spread pulsone. If, in addition, the channel satisfies the crystallization conditions with respect to the period lattice Λ_p, then in an OTFS frame consisting of a spread pilot pulsone and point data pulsones, after cancelling the received signal corresponding to the spread pulsone, we can recover the channel response to any data pulsone.

This translates integration of communication and sensing within a single OTFS frame into geometric properties of a lattice Λ_p used for data transmission and a rotated lattice Λ^* used for sensing. The spread pilot pulsone looks like noise to the point data pulsones, and it is this incoherence that makes it possible to integrate communications and sensing without time-sharing delay-Doppler resources. We demonstrate that integrated sensing and communication increases effective throughput.

Reporter: Mariia Seleznova

Participants

Dr. El Mehdi Achour
Chair of Mathematics of Information Processing
RWTH Aachen
Pontdriesch 12-14
52062 Aachen
GERMANY

Prof. Dr. Rima Alaifari
Departement Mathematik
ETH Zürich
HG G 59.2
Rämistrasse 101
8092 Zürich
SWITZERLAND

Dr. Ernesto Araya
Mathematisches Institut
Ludwig-Maximilians-Universität München
Theresienstr. 39
80333 München
GERMANY

Prof. Dr. Radu Balan
Department of Mathematics
University of Maryland
College Park, MD 20742-4015
UNITED STATES

Prof. Dr. Afonso S. Bandeira
Departement Mathematik
ETH Zürich
Rämistr. 101
8092 Zürich
SWITZERLAND

Wiebke Bartolomaeus
Mathematisches Institut
Ludwig-Maximilians-Universität München
Theresienstr. 39
80333 München
GERMANY

Prof. Dr. Helmut Bölcskei
Mathematical Information Science
ETH Zürich
Room: ETF E 122
Sternwartstrasse 7
8092 Zürich
SWITZERLAND

Dr. Claire Boyer
Laboratoire de Probabilités, Statistique et Modélisation (LPSM), BP 158
Sorbonne Université
Campus Pierre et Marie Curie
4, place Jussieu
75252 Paris Cedex 05
FRANCE

Prof. Dr. A. Robert Calderbank
Gross Hall, Room 317
Pratt School of Engineering
Duke University
140 Science Drive
P.O. Box 90984
Durham, NC 27708
UNITED STATES

Dr. Hung-Hsu Chou
Zentrum Mathematik
TU München
Boltzmannstr. 3
85748 Garching bei München
GERMANY

Cristina Cipriani
Zentrum Mathematik
TU München
Boltzmannstr. 3
85748 Garching bei München
GERMANY

Prof. Dr. Ingrid Daubechies
Department of Mathematics
Duke University
P.O. Box 90984
Durham, NC 27708
UNITED STATES

Prof. Dr. Christine De Mol
Department of Mathematics
Université Libre de Bruxelles
CP 217 Campus Plaine
Boulevard du Triomphe
1050 Bruxelles
BELGIUM

Dr. Sjoerd Dirksen
Mathematical Institute, Utrecht
University
Budapestlaan 6
P.O. Box 80010
3508 TA Utrecht
NETHERLANDS

Dr. Monika Dörfler
Fakultät für Mathematik
Universität Wien
Oskar Morgenstern Platz 1
1090 Wien
AUSTRIA

Dr. Markus Faulhuber
Fakultät für Mathematik
Universität Wien
Oskar-Morgenstern-Platz 1
1090 Wien
AUSTRIA

Adalbert Fono
Mathematisches Institut
Ludwig-Maximilians-Universität
München
Theresienstr. 39
80333 München
GERMANY

Prof. Dr. Simon Foucart
Department of Mathematics
Texas A & M University
College Station, TX 77843-3368
UNITED STATES

Prof. Dr. Remi Gribonval
Laboratoire de l'Informatique du
Parallélisme
ENS de Lyon et INRIA
46 Allée d'Italie
69007 Lyon Cedex
FRANCE

Prof. Dr. Karlheinz Gröchenig
Fakultät für Mathematik
Universität Wien
Oskar-Morgenstern-Platz 1
1090 Wien
AUSTRIA

Prof. Dr. Philipp Grohs
Fakultät für Mathematik
Universität Wien
Oskar Morgenstern Platz 1
1090 Wien
AUSTRIA

Frederik Hoppe
Lehrstuhl für Mathematik der
Informationsverarbeitung
RWTH Aachen University
Pontdriesch 10
52062 Aachen
GERMANY

Prof. Dr. Kathlén Kohn
Department of Mathematics
KTH
Lindstedtsvägen 25
10044 Stockholm
SWEDEN

Christian Koke
Zentrum Mathematik
TU München
Boltzmannstr. 3
85748 Garching bei München
GERMANY

Prof. Dr. Felix Krahmer
Technische Universität München
Department of Mathematics
Boltzmannstraße 3
85748 Garching bei München
GERMANY

Prof. Dr. Richard Küng
Department of Computer Science
Johannes Kepler University Linz
SCP3 0405
Altenberger Strasse 69
4040 Linz
AUSTRIA

Prof. Dr. Gitta Kutyniok
Ludwig-Maximilians-Universität München
Lehrstuhl für Mathematische Grundlagen von Künstlicher Intelligenz
Akademiestraße 7
80799 München
GERMANY

Prof. Dr. Sophie Langer
Department of Applied Mathematics
University of Twente
P.O. Box 217
7500 AE Enschede
NETHERLANDS

Hannah Laus
Zentrum Mathematik
TU München
Boltzmannstr. 3
85748 Garching bei München
GERMANY

Prof. Dr. Johannes Maly
LMU Munich
Akademiestraße 7
80799 München
GERMANY

Maria Matveev
Chair for Mathematical Foundations of Artificial Intelligence
Munich Center for Machine Learning
Ludwig Maximilian University Munich
Akademiestraße 7
80799 München
GERMANY

Dr. Dustin G. Mixon
Department of Mathematics
The Ohio State University
100 Mathematics Building
231 West 18th Avenue
Columbus, OH 43210-1174
UNITED STATES

Prof. Dr. Deanna Needell
Department of Mathematics
University of California at Los Angeles
Los Angeles, CA 90095-1555
UNITED STATES

Prof. Dr. Götz Pfander
Mathematisches Institut für Maschinelles Lernen und Data Science
Mathematisch-Geographische Fakultät
Katholische Universität Eichstätt-Ingolstadt
Hohe-Schul-Straße 5
85049 Ingolstadt
GERMANY

Prof. Dr. Holger Rauhut
Mathematisches Institut
LMU München
Theresienstr. 39
80333 München
GERMANY

Noam Razin
School of Computer Science
Tel Aviv University
Ramat Aviv, Tel Aviv 69978
ISRAEL

Prof. Dr. Karin Schnass
Institut für Mathematik
Universität Innsbruck
Technikerstrasse 13
6020 Innsbruck
AUSTRIA

Prof. Dr. Carola-Bibiane Schönlieb
Department of Applied Mathematics and
Theoretical Physics (DAMTP)
Centre for Mathematical Sciences
Wilberforce Road
Cambridge CB3 0WA
UNITED KINGDOM

Mariia Seleznova
Mathematisches Institut
Ludwig-Maximilians-Universität
München
Akademiestr. 7
80799 München
GERMANY

Anna Shalova
Department of Mathematics
Eindhoven University of Technology
P.O. Box 513
5600 MB Eindhoven
NETHERLANDS

Manjot Singh
Mathematisches Institut
Ludwig-Maximilians-Universität
München
Akademiestr. 7
80799 München
GERMANY

Dr. Mahdi Soltanolkotabi
University of Southern California
3740 McClintock Ave
Los Angeles, CA 95616-8633
UNITED STATES

Prof. Dr. Stefan Steinerberger
Department of Mathematics
University of Washington, Seattle
Seattle, WA 98195
UNITED STATES

Prof. Dr. Dominik Stöger
Mathematisch-Geographische Fakultät
Katholische Universität Eichstätt
Ostenstraße 26-28
85072 Eichstätt
GERMANY

Dr. Laura Thesing
Mathematisches Institut
Ludwig-Maximilians-Universität
München
Akademiestraße 7
80799 München
GERMANY

Dr. Philipp Trunschke
Laboratoire de Mathématiques
Jean Leray UMR 6629
Université de Nantes, B.P. 92208
2 Rue de la Houssiniere
44322 Nantes Cedex 03
FRANCE

Prof. Dr. Felix Voigtlaender
Katholische Universität
Eichstätt-Ingolstadt
Mathematisches Institut für maschinelles
Lernen und Data Science (MIDS)
Auf der Schanz 49
85049 Ingolstadt
GERMANY