Introduction

One of the major advances of science in the 20th century was the discovery of a mathematical formulation of quantum mechanics by Heisenberg in 1925 [103].[1] From a mathematical point of view, this transition from classical mechanics to quantum mechanics amounts to, among other things, passing from the *commutative algebra* of *classical observables* to the *noncommutative algebra* of *quantum mechanical observables*. To understand this better we recall that in classical mechanics an observable of a system (e.g. energy, position, momentum, etc.) is a function on a manifold called the phase space of the system. Classical observables can therefore be multiplied in a pointwise manner and this multiplication is obviously commutative. Immediately after Heisenberg's work, ensuing papers by Dirac [73] and Born–Heisenberg–Jordan [16], made it clear that a quantum mechanical observable is a (selfadjoint) linear operator on a Hilbert space, called the state space of the system. These operators can again be multiplied with composition as their multiplication, but this operation is not necessarily commutative any longer.[2] In fact Heisenberg's *commutation relation*

$$pq - qp = \frac{h}{2\pi i} 1$$

shows that position and momentum operators do not commute and this in turn can be shown to be responsible for the celebrated *uncertainty principle* of Heisenberg. Thus, to get a more accurate description of nature one is more or less forced to replace the commutative algebra of functions on a space by the noncommutative algebra of operators on a Hilbert space.

A little more than fifty years after these developments Alain Connes realized that a similar procedure can in fact be applied to areas of mathematics where the classical notions of space (e.g. measure space, locally compact space, or smooth space) lose its applicability and relevance [37], [35], [36], [39]. The inadequacy of the classical notion of space manifests itself for example when one deals with highly singular "*bad quotients*": spaces such as the quotient of a nice space by the ergodic action of a group, or the space of leaves of a foliation in the generic case, to give just two examples. In all these examples the quotient space is typically ill-behaved, even as a topological space. For instance it may fail to be even Hausdorff, or have enough open sets, let alone being a reasonably smooth space. The unitary dual of a discrete group, except when the group is abelian or almost abelian, is another example of an ill-behaved space.

[1] A rival proposal which, by the Stone–von Neumann uniqueness theorem, turned out to be essentially equivalent to Heisenberg's was arrived at shortly afterwards by Schrödinger [173]. It is however Heisenberg's *matrix mechanics* that directly and most naturally relates to noncommutative geometry.

[2] Strictly speaking selfadjoint operators do not form an algebra since they are not closed under multiplication. By an algebra of observables we therefore mean the algebra that they generate.

One of Connes' key observations is that in all these situations one can define a noncommutative algebra through a universal method which we call the *noncommutative quotient construction* that captures most of the information hidden in these unwieldy quotients. Examples of this noncommutative quotient construction include the crossed product by an action of a group, or in general by an action of a groupoid. In general the noncommutative quotient is the groupoid algebra of a topological groupoid.

This new notion of geometry, which is generally known as *noncommutative geometry*, is a rapidly growing new area of mathematics that interacts with and contributes to many disciplines in mathematics and physics. Examples of such interactions and contributions include: the theory of operator algebras, index theory of elliptic operators, algebraic and differential topology, number theory, the Standard Model of elementary particles, the quantum Hall effect, renormalization in quantum field theory, and string theory.

To understand the basic ideas of noncommutative geometry one should perhaps first come to grips with the idea of a *noncommutative space*. What is a noncommutative space? The answer to this question is based on one of the most profound ideas in mathematics, namely a *duality* or *correspondence* between algebra and geometry,[3]

$$\boxed{\text{Algebra} \longleftrightarrow \text{Geometry}}$$

according to which every concept or statement in Algebra corresponds to, and can be equally formulated by, a similar concept and statement in Geometry.

On a physiological level this correspondence is perhaps related to a division in the human brain: one computes and manipulates symbols with the left hemisphere of the brain and one visualizes things with the right hemisphere. Computations evolve in time and have a temporal character, while visualization is instant and immediate. It was for a good reason that Hamilton, one of the creators of modern algebraic methods, called his approach to algebra, e.g. to complex numbers and quaternions, the *science of pure time* [101].

We emphasize that the algebra-geometry correspondence is by no means a new observation or a new trend in mathematics. On the contrary, this duality has always existed and has been utilized in mathematics and its applications very often. The earliest example is perhaps the use of numbers in counting. It is, however, the case that throughout history each new generation of mathematicians has found new ways of formulating this principle and at the same time broadening its scope. Just to mention a few highlights of this rich history we quote Descartes (analytic geometry), Hilbert (affine varieties and commutative algebras), Gelfand–Naimark (locally compact spaces and commutative C^*-algebras), and Grothendieck (affine schemes and commutative

[3] For a modern and very broad point of view on this duality, close to the one adopted in this book, read the first section of Shafarevich's book [176] as well as Cartier's article [31].

rings). A key idea here is the well-known relation between a space and the commutative algebra of functions on that space. More precisely, there is a duality between certain categories of geometric spaces and the corresponding categories of algebras representing those spaces. Noncommutative geometry builds on, and vastly extends, this fundamental duality between classical geometry and commutative algebras.

For example, by a celebrated theorem of Gelfand and Naimark [91], one knows that the information about a compact Hausdorff space is fully encoded in the algebra of continuous complex-valued functions on that space. The space itself can be recovered as the space of maximal ideals of the algebra. Algebras that appear in this way are commutative C^*-algebras. This is a remarkable theorem since it tells us that any natural construction that involves compact spaces and continuous maps between them has a purely algebraic reformulation, and vice-versa any statement about commutative C^*-algebras and C^*-algebraic maps between them has a purely geometric-topological meaning.

Thus one can think of the category of not necessarily commutative C^*-algebras as the dual of an, otherwise undefined, category of *noncommutative locally compact spaces*. What makes this a successful proposal is, first of all, a rich supply of examples and, secondly, the possibility of extending many of the topological and geometric invariants to this new class of 'spaces' and applications thereof.

Noncommutative geometry has as its special case, in fact as its limiting case, classical geometry, but geometry expressed in algebraic terms. In some respect this should be compared with the celebrated *correspondence principle* in quantum mechanics where classical mechanics appears as a limit of quantum mechanics for large quantum numbers or small values of Planck's constant. Before describing the tools needed to study noncommutative spaces let us first briefly recall a couple of other examples from a long list of results in mathematics that put in duality certain categories of geometric objects with a corresponding category of algebraic objects.

To wit, Hilbert's Nullstellensatz states that the category of affine algebraic varieties over an algebraically closed field is equivalent to the opposite of the category of finitely generated commutative algebras without nilpotent elements (so-called reduced algebras). This is a perfect analogue of the Gelfand–Naimark theorem in the world of algebraic geometry. Similarly, Swan's (resp. Serre's) theorem states that the category of vector bundles over a compact Hausdorff space (resp. over an affine algebraic variety) X is equivalent to the category of finitely generated projective modules over the algebra of continuous functions (resp. the algebra of regular functions) on X.

A pervasive idea in noncommutative geometry is to treat certain classes of noncommutative algebras as noncommutative spaces and to try to extend tools of geometry, topology, and analysis to this new setting. It should be emphasized, however, that, as a rule, this extension is hardly straightforward and most of the times involves surprises and new phenomena. For example, the theory of the flow of weights and the corresponding modular automorphism group in von Neumann algebras [41] has no counterpart in classical measure theory, though the theory of von Neumann algebras is generally regarded as noncommutative measure theory. Similarly, as we shall see in

Chapters 3 and 4 of this book, the extension of de Rham (co)homology of manifolds to cyclic (co)homology for noncommutative algebras was not straightforward and needed some highly non-trivial considerations. As a matter of fact, de Rham cohomology can be defined in an algebraic way and therefore can be extended to all commutative algebras and to all schemes. This extension, however, heavily depends on exterior products of the module of Kähler differentials and on the fact that one works with commutative algebras. In the remainder of this introduction we focus on topological invariants that have proved very useful in noncommutative geometry.

Of all topological invariants for spaces, topological K-theory has the most straightforward extension to the noncommutative realm. Recall that topological K-theory classifies vector bundles on a topological space. Motivated by the above-mentioned Serre–Swan theorem, it is natural to define, for a not necessarily commutative ring A, $K_0(A)$ as the group defined by the semigroup of isomorphism classes of finite projective A-modules. Provided that A is a Banach algebra, the definition of $K_1(A)$ follows the same pattern as for spaces, and the main theorem of topological K-theory, the Bott periodicity theorem, extends to all Banach algebras [14].

The situation was much less clear for K-homology, a dual of K-theory. By the work of Atiyah [6], Brown–Douglas–Fillmore [22], and Kasparov [115], one can say, roughly speaking, that K-homology cycles on a space X are represented by abstract elliptic operators on X and, whereas K-theory classifies the vector bundles on X, K-homology classifies the abstract elliptic operators on X. The pairing between K-theory and K-homology takes the form $\langle [D], [E] \rangle = \text{index}(D_E)$, the Fredholm index of the elliptic operator D with coefficients in the 'vector bundle' E. Now one good thing about this way of formulating K-homology is that it almost immediately extends to noncommutative C^*-algebras. The two theories are unified in a single theory called KK-theory, due to Kasparov [115].

Cyclic cohomology was discovered by Connes in 1981 [36], [39] as the right noncommutative analogue of the de Rham homology of currents and as a receptacle for a noncommutative Chern character map from K-theory and K-homology. One of the main motivations was transverse index theory on foliated spaces. Cyclic cohomology can be used to identify the K-theoretic index of transversally elliptic operators which lie in the K-theory of the noncommutative algebra of the foliation. The formalism of cyclic cohomology and noncommutative Chern character maps form an indispensable part of noncommutative geometry. A very interesting recent development in cyclic cohomology theory is the *Hopf cyclic cohomology* of Hopf algebras and Hopf (co)module (co)algebras in general. Motivated by the original work in [57], [58] this theory has now been extended in [98], [99].

The following "dictionary" illustrates noncommutative analogues of some of the classical theories and concepts originally conceived for spaces. In this book we deal only with a few items of this ever expanding dictionary.

commutative	**noncommutative**
measure space	von Neumann algebra
locally compact space	C^*-algebra
vector bundle	finite projective module
complex variable	operator on a Hilbert space
infinitesimal	compact operator
range of a function	spectrum of an operator
K-theory	K-theory
vector field	derivation
integral	trace
closed de Rham current	cyclic cocycle
de Rham complex	Hochschild homology
de Rham cohomology	cyclic homology
Chern character	Connes–Chern character
Chern–Weil theory	noncommutative Chern–Weil theory
elliptic operator	K-cycle
spin^c Riemannian manifold	spectral triple
index theorem	local index formula
group, Lie algebra	Hopf algebra, quantum group
symmetry	action of Hopf algebra

Noncommutative geometry is already a vast subject. This book is an introduction to some of its basic concepts suitable for graduate students in mathematics and physics. While the idea was to write a primer for the novice to the subject, some acquaintance with functional analysis, differential geometry and algebraic topology at a first year graduate level is assumed. To get a better sense of the beauty and depth of the subject the reader can go to no better place than the authoritative book [41]. There are also several introductions to the subject, with varying lengths and attention to details, that the reader can benefit from [94], [186], [155], [106], [129], [119], [144], [146], [70], [69]. They each emphasize rather different aspects of noncommutative geometry. For the most complete account of what has happened in the subject after the publication of [41], the reader should consult [55] and references therein.

To summarize our introduction we emphasize that what makes the whole project of noncommutative geometry a viable and extremely important enterprize are the following three fundamental points:

• There is a vast repertoire of noncommutative spaces and there are very general methods to construct them. For example, consider a *bad quotient* of a nice and smooth space by an equivalence relation. Typically the (naive) quotient space is not even Hausdorff and has very bad singularities, so that it is beyond the reach of classical geometry and topology. Orbit spaces of group actions and the space of leaves of a foliation are examples of this situation. In algebraic topology one replaces such naive quotients by homotopy quotients, by using the general idea of a classifying space. This is however not good enough and not general enough, as the classifying space is only a homotopy construction and does not see any of the smooth structure. A key observation throughout [41] is that in all these situations one can attach a noncommutative space, e.g. a (dense subalgebra of a) C^*-algebra or a von Neumann algebra, that captures most of the information hidden in these quotients. The general construction starts by first replacing the equivalence relation by a groupoid and then considering the associated groupoid algebra in its various completions. We shall discuss this technique in detail in Chapter 2 of this book.

• The possibility of extending many of the tools of classical geometry and topology that are used to probe classical spaces to this noncommutative realm. The topological K-theory of Atiyah and Hirzebruch, and its dual theory known as K-homology, as well as the Bott periodicity theorem, have a natural extension to the noncommutative world [14]. Finding the right noncommutative analogue of de Rham cohomology and Chern–Weil theory was less obvious and was achieved thanks to the discovery of cyclic cohomology [36], [38]. In Chapters 3 and 4 of this book we shall give a detailed account of cyclic cohomology and its relation with K-theory and K-homology. Another big result of recent years is the local index formula of Connes and Moscovici [57]. Though we shall not discuss it in this book, it suffices to say that this result comprises a vast extension of the classical Atiyah–Singer index theorem to the noncommutative setup.

• *Applications*. Even if we wanted to restrict ourselves just to classical spaces, methods of noncommutative geometry would still be very relevant and useful. For example, a very natural and general proof of the Novikov conjecture on the homotopy invariance of higher signatures of non-simply connected manifolds (with word hyperbolic fundamental groups) can be given using the machinery of noncommutative geometry [56]. The relevant noncommutative space here is the (completion of the) group ring of the fundamental group of the manifold. We also mention the geometrization of the Glashow–Weinberg–Salam Standard Model of elementary particles via noncommutative geometry (cf. [55] and references therein). Moving to more recent applications, we mention the approach to the Riemann hypothesis and the spectral realization of zeros of the zeta function via noncommutative spaces [18], [42] as well as the mathematical underpinning of renormalization in quantum field theory as a Riemann–Hilbert Correspondence [50], [51]. These results have brought noncommutative geometry much closer to central areas of modern number theory, algebraic geometry and high energy physics. We shall not follow these developments in this book. For a complete and up to date account see [55].

Let me now briefly explain the contents of this book. Chapter 1 describes some of the fundamental algebra-geometry correspondences that are vital for a better understanding of noncommutative geometry. The most basic ones, for noncommutative geometry at least, are the Gelfand–Naimark and the Serre–Swan theorems. They lead to ideas of noncommutative space and noncommutative vector bundles. We give several examples of noncommutative spaces, most notably noncommutative tori, group C^*-algebras, and quantum groups. The last section of this first chapter is a self contained introduction to Hopf algebras and quantum groups and the idea of symmetry in noncommutative geometry. Chapter 2 is about forming noncommutative quotients via groupoids and groupoid algebras. This is one of the most universal and widely used methods for constructing noncommutative spaces. Another important concept in this chapter is the idea of Morita equivalence of algebras, both at purely algebraic and C^*-algebraic levels. Among other things, Morita equivalence clarifies the relation between noncommutative quotients and classical quotients. Chapter 3 is devoted to cyclic (co)homology, its relation with Hochschild (co)homology through Connes' long sequence and spectral sequence, and its relation with de Rham (co)homology. Three different definitions of cyclic cohomology are given in this chapter, each shedding light on a different aspect of the theory. Continuous versions of cyclic and Hochschild theory for topological algebras is developed in this section. This plays an important role in applications. We also give several important examples of algebras for which these invariants are fully computed. In Chapter 4 we define the Connes–Chern character map for both K-theory and K-homology. For K-theory it is the noncommutative analogue of the classical Chern character map from K-theory to de Rham cohomology. It can also be described as a pairing between K-theory and cyclic cohomology. Fredholm modules, as cycles for K-homology, are introduced next and, for finitely summable Fredholm modules, their Connes–Chern character with values in cyclic cohomology is introduced. Then we use these pairings to prove an index formula from [39] relating the analytic Fredholm index of a finitely summable Fredholm module to its topological index. This is an example of an index formula in noncommutative geometry. The very last section of this chapter summarizes many ideas of the book into one commuting diagram which is the above mentioned index formula. In an effort to make this book as self-contained as possible, we have added four appendices covering basic material on C^*-algebras, compact and Fredholm operators, projective modules, and some basic category theory language.

This book is partly based on series of lectures I gave at the Fields Institute in Toronto, Canada, in Fall 2002 and at the second and fifth Annual Spring Institute on Noncommutative Geometry and Operator Algebras in spring 2004 and spring 2007 at Vanderbilt University, USA.

It is a great pleasure to acknowledge my deepest gratitude to Alain Connes for his invaluable advice, support, and encouragement over many years and during the writing of this book. I would like to warmly thank Klaas Landsman and Joseph Varilly who carefully read the entire manuscript and provided me with many valuable suggestions and corrections. I would also like to sincerely thank Manfred Karbe and

Irene Zimmermann of the European Mathematical Society Publishing House for taking care of the text and technical production, Matilde Marcolli for encouraging me to publish the book at an early stage, and Katia Consani for inviting me to give a series of lectures at Johns Hopkins University in March 2008 based on a pre-final version of this book. I also thank Piotr Hajac and Gianni Landi for reading the text and for their valuable suggestions. Finally I would like to thank Arthur Greenspoon for carefully proofreading the whole text.