Introduction

J. Toby Stafford

There are multiple interactions between noncommutative algebra and representation theory on the one hand and classical algebraic geometry on the other, and the aim of this book is to expand upon this interplay. One of the most obvious areas of interaction is in noncommutative algebraic geometry, where the ideas and techniques of algebraic geometry are used to study noncommutative algebra. An introduction to this material is given in Chapter I. Many of the algebras that appear naturally in that, and other, areas of mathematics are deformations of commutative algebras, and so in Chapter II we provide a comprehensive introduction to that theory. One of the most interesting classes of algebras to have appeared recently in representation theory, and discussed in Chapter III, is that of symplectic reflection algebras. Finally, one of the strengths of these topics is that they have applications back in the commutative universe. Illustrations of this appear throughout the book, but one particularly important instance is that of noncommutative (crepant) resolutions of singularities. This forms the subject of Chapter IV.

These notes have been written up as an introduction to these topics, suitable for advanced graduate students or early postdocs. In keeping with the lectures upon which the book is based, we have included a large number of exercises, for which we have given partial solutions at the end of book. Some of these exercises involve computer computations, and for these we have either included the code or indicated web sources for that code.

We now turn to the individual topics in this book. Throughout the introduction k will denote an algebraically closed base field and all algebras will be k-algebras.

I. Noncommutative projective geometry. This subject seeks to use the results and intuition from algebraic geometry to understand noncommutative algebras. There are many different versions of noncommutative algebraic geometry, but the one that concerns us is noncommutative projective algebraic geometry, as introduced by Artin, Tate, and Van den Bergh [9, 10].

As is true of classical projective algebraic geometry, we will be concerned with *connected graded (cg) k-algebras* A. This means that (1) $A = \bigoplus_{n\geq 0} A_n$ with $A_n A_m \subseteq A_{n+m}$ for all $n, m \geq 0$ and (2) $A_0 = k$. For the rest of the introduction we will also assume, for simplicity, that A is generated as a k-algebra by the finite-dimensional vector space A_1. Write $A_+ = \bigoplus_{n>0} A_n$ for the irrelevant ideal.

The starting point to this theory appears in work of Artin and Schelter [5], who were interested in classifying noncommutative analogues A of the polynomial ring $k[x, y, z]$ or, as we will describe later, noncommutative analogues of $\mathbb{P}^2$. So, what should the definition be? The first basic condition is that A should have finite global dimension gldim $A = m$, in the sense that every finitely generated A-module M should have a finite projective resolution. This hypothesis is insufficient by itself; for example the free algebra $k\{x, y\}$ has global dimension one. So we also demand that A have *polynomially bounded growth* in the sense that the function $p(n) = \dim_k A_n$ is bounded above by some polynomial function of n. This is still not enough to eliminate rings like $k\{x, y\}/(xy)$ that have rather unpleasant properties. The insight of Artin and Schelter was to add a *Gorenstein* condition: $\mathrm{Ext}^i_A(k, A) = \delta_{i,m}k$, where k is the trivial (right) A-module A/A_+. In the commutative case this condition is equivalent to the ring having finite injective dimension, hence weaker than having finite global dimension, yet in many ways in the noncommutative setting it is a more stringent condition. Algebras with these three properties — global dimension m, polynomially bounded growth and the Gorenstein condition — are now called *Artin–Schelter regular* or *AS-regular rings of dimension m*. These algebras appear throughout noncommutative algebraic geometry and form the underlying theme for Chapter I. All references in this subsection are to that chapter.

Artin–Schelter regular algebras of dimension 2 are easily classified; this is the content of Theorem 2.2.1. In fact there are just two examples: the *quantum plane* $k_q[x, y] := k\{x, y\}/(xy - qyx)$ for $q \in k \smallsetminus \{0\}$ and the *Jordan plane* $k_J[x, y] := k\{x, y\}/(xy - yx - y^2)$. (Since we are concerned with projective rather than affine geometry, we probably ought to call them the quantum and Jordan projective lines, but we will stick to these more familiar names.) It is straightforward to analyse the properties of these rings using elementary methods.

So it was the case of dimension 3 that interested Artin and Schelter, and here things are not so simple. The Gorenstein condition enables one to obtain detailed information about the projective resolution of the trivial module $k = A/A_+$. In many cases this is enough to describe the algebra in considerable detail, and in particular to give a basis for the algebra. However there was one algebra, now called the *Sklyanin algebra*, that Artin and Schelter could not completely understand (this algebra is described in terms of generators and relations in Example 1.3.4 but its precise description is not so important here). It was the elucidation of this and closely related algebras that required the introduction of geometric techniques through the work of Artin, Tate, and Van den Bergh [9, 10].

The idea is as follows. Given a commutative cg domain A the (closed) points of the corresponding projective variety Proj(A) can be identified with the maximal nonirrelevant graded prime ideals; under our hypotheses these are the graded ideals P such that $A/P \cong k[x]$ is a polynomial ring in one variable. In the noncommutative case, this is too restrictive — for example if q is not a root of unity, then $k_q[x, y]$ has just two such ideals; (x) and (y). Instead one works module-theoretically and defines a *point module* to be a right A-module $M = \bigoplus_{n \geq 0} M_n$ such that $M = M_0 A$ and $\dim_k M_n = 1$ for all $n \geq 0$. Of course, when A is commutative, these are the factor rings we just mentioned, but they are more subtle in the noncommutative setting and are discussed in detail in Sections 3 and 4. The gist is as follows. Let A be an AS regular algebra of dimension 3. Then the point modules for A are in one-to-one correspondence with (indeed, parametrised by) a scheme E, known naturally enough as the *point scheme* of A. This scheme further comes equipped with the extra data of an automorphism σ and a line bundle $\mathcal{L}$. From these data one can construct an algebra, known as the *twisted homogeneous coordinate ring* $B = B(E, \mathcal{L}, \sigma)$ of E. If $A = k[x_0, x_1, x_2]$ were a commutative polynomial ring in three variables, then $E = \mathbb{P}^2$ and B would simply be A. In the noncommutative case E will either be a surface (indeed either $\mathbb{P}^2$ or $\mathbb{P}^1 \times \mathbb{P}^1$) in which case $A = B$ or, more interestingly, E could be a curve inside one of those surfaces. The interesting case is when E is an elliptic curve, as is the case for the Sklyanin algebra we mentioned before. This also helps explain why the Sklyanin algebra caused such a problem in the original work of Artin and Schelter: elliptic curves are not so easily approached by the sorts of essentially linear calculations that were integral to their work.

The beauty of this theory is that the geometry of the point scheme E can be used to describe the twisted homogeneous coordinate ring $B = B(E, \mathcal{L}, \sigma)$ and its modules in great detail. Moreover, for an AS regular algebra A of dimension 3, the ring B is a factor $B = A/gA$ of A, and the pleasant properties of B lift to give a detailed description of A and ultimately to classify the AS-regular algebras of dimension 3. This process is outlined in Section 3.2. An important and surprising consequence is that these algebras A are all noetherian domains; thus every right (or left) ideal of A is finitely generated.

We study twisted homogeneous coordinate rings in some detail since they are one of the basic notions in the subject, with numerous applications. A number of these applications are given in Section 5. For example, if A is a domain for which $\dim_k A_n$ grows linearly, then, up to a finite-dimensional vector space, A is a twisted homogeneous coordinate ring (see Theorem 5.1.1 for the details). One consequence of this is that the module structure of the algebra A is essentially that of a commutative ring. To explain the module theory we need some more notation.

If A is a commutative cg algebra then one ignores the irrelevant ideal A_+ in constructing the projective variety $\mathrm{Proj}(A)$. This means we should ignore finite-dimensional modules when relating that geometry to the module structure of A. This holds in the noncommutative case as well. Assume that A is noetherian, which is the case that interests us, and let $\mathrm{gr}\, A$ denote the category of finitely generated graded A-modules $M = \bigoplus_{n \in \mathbb{Z}} M_i$ (thus $M_i A_j \subseteq M_{i+j}$ for all i and j). The category $\mathrm{qgr}(A)$ is defined to be the quotient category of $\mathrm{gr}\, A$ by the finite-dimensional modules; see Definition 4.0.7 for more details. A surprisingly powerful intuition is to regard $\mathrm{qgr}(A)$ as the category of coherent sheaves on the (nonexistent) space $\mathrm{Proj}(A)$. Similarly, there are strong arguments for saying that the AS regular algebras A of dimension 3 (or at least those for which $\dim_k A_1 = 3$) are the coordinate rings of the noncommutative $\mathbb{P}^2$'s.

The fundamental result relating algebraic geometry to twisted homogeneous coordinate rings is the following theorem of Artin and Van den Bergh [11], which is itself a generalisation of a result of Serre [198]: under a condition called σ-ampleness of the sheaf $\mathcal{L}$, the category qgr $B(E, \mathcal{L}, \sigma)$ is equivalent to the category of coherent sheaves on E. In particular, if A is a domain for which $\dim_k A_n$ grows linearly with n, as was the case two paragraphs ago, then qgr A will be equivalent to the category of coherent sheaves on a projective curve: sometimes this is phrased as saying that noncommutative curves are commutative!

Structure of Chapter I. The main aim of this chapter is to give the reader a firm understanding of the mathematics behind the above outline, and we have kept the geometric prerequisites to a minimum. Thus, in Section 1 we emphasise techniques for calculating the basis (or more generally the Hilbert series) of a graded algebra given by generators and relations. Section 2 introduces the Artin–Schelter regular algebras and, again, we emphasise how to use the Gorenstein condition to understand some of the basic examples. Of course this does not work everywhere, so Section 3 introduces point modules, the corresponding point scheme and shows how to compute this in explicit examples. Section 4 then describes the corresponding twisted homogeneous coordinate rings, while Section 5 outlines the applications of these techniques to the classifications of noncommutative curves and particular classes of noncommutative surfaces.

II. Deformations of algebras in noncommutative geometry. For simplicity, in discussing Chapter II we will assume that the base field k has characteristic zero. A great many algebras appearing in noncommutative algebra, and certainly most of the ones described in this book, are deformations of commutative algebras. For example, if $k_q[x, y] = k\{x, y\}/(xy - qyx)$ is the quantum plane mentioned above then it is easy to see that this algebra has basis $\{x^i y^j\}$. Thus, as q passes from 1 to a general element of k, it is natural to regard this algebra as *deforming* the multiplication of the algebra $k[x, y]$. In fact there are many different ways of deforming algebras and some very deep results about when this is possible. This is the topic of Chapter II. Once again, all references in this subsection are to that chapter.

Here are a couple of illustrative examples. Given a finite-dimensional Lie algebra $\mathfrak{g}$ over the field k, with Lie bracket $\{-, -\}$, its enveloping algebra $U\mathfrak{g}$ is defined to be the factor of the tensor algebra $T\mathfrak{g}$ on $\mathfrak{g}$ modulo the relations $xy - yx - \{x, y\}$ for $x, y \in \mathfrak{g}$. One can also form the symmetric algebra $\operatorname{Sym}\mathfrak{g}$ on $\mathfrak{g}$, which is nothing more than the polynomial ring in $\dim_k \mathfrak{g}$ variables. Perhaps the most basic theorem on enveloping algebras is the PBW or Poincaré–Birkhoff–Witt Theorem: if one filters $U\mathfrak{g} = \bigcup_{n \geq 0} \Lambda_{\leq n}$ by assigning $\mathfrak{g} + k$ to $\Lambda_{\leq 1}$, then $\operatorname{Sym}\mathfrak{g}$ is isomorphic to the associated graded ring $\operatorname{gr} U\mathfrak{g} = \bigoplus \Lambda_{\leq n}/\Lambda_{\leq (n-1)}$. We interpret this as saying that $U\mathfrak{g}$ *is a filtered deformation of* $\operatorname{Sym}\mathfrak{g}$. A similar phenomenon occurs with the *Weyl algebra*, or ring of linear differential operators on $\mathbb{C}^n$. This is the ring with generators $\{x_i, \partial_i : 1 \leq i \leq n\}$ with relations $\partial_i x_i - x_i \partial_i = 1$ and all other generators commuting. Again one can filter this algebra by putting the x_i and ∂_j into degree one and its associated graded ring is then the polynomial ring $\mathbb{C}[x_1, \ldots, y_n]$ in $2n$ variables. As is indicated in Section 1, there are numerous other examples of filtered deformations of commutative algebras, including more general rings of differential operators and even some of the algebras from Chapter III.

The commutative rings B that arise as the associated graded rings $B = \operatorname{gr} A = \bigoplus \Lambda_{\leq n}/\Lambda_{\leq(n-1)}$ of filtered rings $A = \bigcup \Lambda_{\leq n}$ automatically have the extra structure of a Poisson algebra. Indeed, given non-zero elements $\bar{a} \in \Lambda_{\leq n}/\Lambda_{\leq(n-1)}$ and $\bar{b} \in \Lambda_{\leq m}/\Lambda_{\leq(m-1)}$, with preimages $a, b \in A$ then we define a new bracket $\{\bar{a}, \bar{b}\} = ab - ba \mod \Lambda_{\leq(m+n-1)}$. It is routine to see that this is actually a *Poisson bracket* in the sense that it is a Lie bracket satisfying the Leibniz identity $\{ab, c\} = a\{b, c\} + b\{a, c\}$. The algebra $\operatorname{gr} A$ is then called a *Poisson algebra*.

One can ask if the reverse procedure holds: Given a commutative Poisson algebra B, can one deform it to a noncommutative algebra A in such a way that the Poisson structure on B is induced from the multiplication in A? This is better phrased in terms of infinitesimal and formal deformations, but see Corollary 2.6.6 for the connection. To describe these deformations, pick an augmented base commutative ring R with augmentation ideal R_+, which for us means either $R = k[[h]]$ or $R = k[h]/(h^n)$ with $R_+ = hR$. Then a *(flat) deformation* of B over R is (up to some technicalities) an R-algebra A, isomorphic to $B \otimes_k R$ as an R-module, such that $A \otimes_R R/R_+ = B$ as k-algebras. In other words, a deformation of B over R is an algebra $(B \otimes_k R, \cdot)$ such that $a \cdot b = ab \mod R_+$. An *infinitesimal deformation* of B is a flat deformation over $R = k[h]/(h^2)$, while a *formal deformation* is the case when $R = k[[h]]$. In both cases the multiplication on A induces a Poisson structure on B by $\{\bar{a}, \bar{b}\} = h^{-1}(ab - ba)$ (which does make sense in the infinitesimal case) and we require that this is the given Poisson structure on B. Remarkably, these concepts are indeed equivalent: *Poisson structures on the coordinate ring B of a smooth affine variety X correspond bijectively to formal deformations of B*. However it takes much more work to make this precise (in particular one needs to work with appropriate equivalence classes on the two sides) and much of Sections 3 and 4 is concerned with setting this up. The original result here is Kontsevich's famous formality theorem, which was first proved at the level of $\mathbb{R}^n$ or more generally C^∞ manifolds. Kontsevich also outlined how to extend this to smooth affine (and some nonaffine) algebraic varieties, while a thorough study in the global algebraic setting was accomplished by Yekutieli and others; see Section 4.6 for the details.

The starting point to deformation theory is that deformations are encoded in Hochschild cohomology. To be a little more precise, let B be a k-algebra with opposite ring B^{op} and set $B^e = B \otimes_k B^{\mathrm{op}}$. The infinitesimal deformations of B are encoded by the second Hochschild cohomology group $HH^2(B) = \operatorname{Ext}^2_{B^e}(B, B)$, while the obstructions to extending these deformations to higher-order ones (i.e., those where $R = k[h]/(h^2)$ is replaced by some $R = k[h]/(h^n)$) are contained within the third Hochschild cohomology group HH^3. This is made precise in Section 3 and put into a more general context in Section 4.

One disadvantage of formal deformation theorems like Kontsevich's is that the multiplication on the deformation is very complicated to describe, yet in many concrete examples (like enveloping algebras) there is actually quite a simple formula for that multiplication. The possible deformations also have a range of different properties, and can often be "smoother" than the original commutative algebra. A illustrative example is given by the fixed ring $\mathbb{C}[x, y]^G$ where the generator σ of $G = \mathbb{Z}/(2)$ acts by -1 on x and y. This has many interesting deformations, including the factor $\overline{U} = U(\mathfrak{sl}_2)/(\Omega)$ of the enveloping algebra of $\mathfrak{sl}_2$ by its Casimir element. (Here $\overline{U}$ is smooth in the sense that, for instance, it has finite global dimension, whereas the global dimension of $\mathbb{C}[x, y]^G$ is infinite.) The ring $\overline{U}$ in turn has many different interpretations; for example, as the ring of global differential operators on the projective line (see Theorems 1.8.2 and (1.J)) or as a spherical subalgebra of a Cherednik algebra in Chapter III.

This example can be further generalised to the notion of a Calabi–Yau algebra. These algebras are ubiquitous in this book. The formal definition is given in Definition 3.7.9 but here we simply note that connected graded Calabi–Yau algebras are a special case of AS regular algebras (see Section 5.5.3 of Chapter I). In particular the polynomial ring $\mathbb{C}[x_1, \dots, x_n]$ is Calabi–Yau, as are many of its deformations, including Weyl algebras and many enveloping algebras. Further examples are provided by $U(\mathfrak{sl}_2)/(\Omega)$ and the symplectic reflection algebras of Chapter III, as well as various noncommutative resolutions of Chapter IV. As these examples suggest, Calabi–Yau algebras can frequently be written as deformations of commutative rings or at least of rings that are "close" to commutative. This is discussed in Section 5 and has important applications to both commutative and noncommutative algebras, as is explained in the next two subsections.

Structure of Chapter II. The aim of the chapter is to give an introduction to deformation theory. Numerous motivating examples appear in Section 1, including enveloping algebras, rings of differential operators and Poisson algebras. The basic concepts of formal deformation theory and Hochschild (co)homology appear in Section 2, while the relationship between these concepts is examined in greater depth in Section 3. These ideas are considerably generalised in Section 4, in order to give the appropriate context for Kontsevich's formality theorem. The ramifications of this result and a hint to its proof are also given there. Finally, Section 5 discusses Calabi–Yau algebras and their applications to deformation theory, such as to quantizations of isolated hypersurface singularities.

III. Symplectic reflection algebras. A fascinating class of algebras that have only recently been discovered (the first serious treatment appears in the seminal paper of Etingof and Ginzburg [99] from 2002) are the *symplectic reflection algebras*, also known in a special case as *rational Cherednik algebras*. They have many interactions with, and applications to, other parts of mathematics and are also related to deformation theory, noncommutative algebraic geometry and noncommutative resolutions. As such, they form a natural class of algebras to study in depth in this book, and we do so in Chapter III. Once again, all references in this subsection are to that chapter.

We first describe these algebras as deformations. Let G be a finite subgroup of $\mathrm{GL}(V)$ for a finite-dimensional vector space V, say over $\mathbb{C}$ for simplicity. Then G acts naturally on the coordinate ring $\mathbb{C}[V]$ and a classic theorem of Chevalley–Shephard–Todd says that the quotient variety $V/G = \operatorname{Spec}\mathbb{C}[V]^G$ is smooth if and only if G is a complex reflection group (see Section 1 for the definitions). There is a symplectic analogue of reflection groups, where V is now symplectic and $G \subset \mathrm{Sp}(V)$. (The simplest case, of type A_{n-1}, is when $G = S_n$ is the symmetric group acting naturally on $\mathbb{C}^n \oplus (\mathbb{C}^n)^*$ by simultaneous permutations of the coordinates.) The variety V/G will not now be smooth; for example in the A_1 case V/G is the surface $xy = z^2$. However there are some very natural noncommutative deformations of $\mathbb{C}[V/G] := \mathbb{C}[V]^G$; notably $\overline{U}_\lambda = U(\mathfrak{sl}_2)/(\Omega - \lambda)$, where Ω is again the Casimir element and $\lambda \in \mathbb{C}$. For all but one choice of λ, the ring $\overline{U}_\lambda$ has finite global dimension and can be regarded as a smooth noncommutative deformation of $\mathbb{C}[V/G]$.

This generalises to any symplectic reflection group. Given such a group $G \subset \mathrm{Sp}(V)$, one can form the invariant ring $\mathbb{C}[V]^G$ and the *skew group ring* $\mathbb{C}[V] \rtimes G$; this is the same abelian group as the ordinary group ring $\mathbb{C}[V]G$, except that the multiplication is twisted: $gf = f^g g$ for $f \in \mathbb{C}[V]$ and $g \in G$. Then Etingof and Ginzburg [99] showed that one can deform $\mathbb{C}[V] \rtimes G$ into a noncommutative algebra, called *the symplectic reflection algebra* $H_{t,c}(G)$, depending on two parameters t and c. The trivial idempotent $e = \sum_{g\in G} g|G|^{-1}$ still lives in this ring and the *spherical subalgebra* $eH_{t,c}(G)e$ is then a deformation of $\mathbb{C}[V]^G$. Crucially, these algebras are filtered deformations in the sense of Chapter II and so, under a natural filtration, one has an analogue of the PBW Theorem: $\operatorname{gr} H_{t,c}(G) = \mathbb{C}[V] \rtimes G$ and $\operatorname{gr} eH_{t,c}(G)e = \mathbb{C}[V]^G$.

The parameter t can always be scaled and so can be chosen to be either 0 or 1. These cases are very different. For most of the chapter we will work in the case $t = 1$ and write $H_c(G) = H_{1,c}(G)$.

The rings $H_c(G)$ are typically defined in terms of generators and relations, which are not easy to unravel (see Definition 1.2.1 and Equation 1.C). However, in the A_1 case, $eH_c(G)e = \overline{U}_\lambda$ for some $\lambda \in \mathbb{C}$, and all such λ occur. In general the properties of the spherical subalgebras $eH_{t,c}(G)e$ are reminiscent of those of a factor ring of an enveloping algebra of a semisimple Lie algebra, and this analogy will guide much of the exposition.

This similarity is most apparent in the special case of Cherednik algebras. Here one takes a complex reflection group $W \subseteq \mathrm{GL}(\mathfrak{h})$ for a complex vector space $\mathfrak{h}$. Then W acts naturally on $V = \mathfrak{h} \times \mathfrak{h}^*$ and defines a symplectic reflection group. The *rational Cherednik algebra* is then the corresponding symplectic reflection algebra $H_{t,c}(W)$. Inside $H_{t,c}(W)$, one has copies of $\mathbb{C}[\mathfrak{h}]$ and $\mathbb{C}[\mathfrak{h}^*]$ as well as the group ring $\mathbb{C}W$ and the PBW Theorem can be refined to give a triangular decomposition $H_{t,c}(G) \cong \mathbb{C}[\mathfrak{h}] \otimes_{\mathbb{C}} \mathbb{C}W \otimes_{\mathbb{C}} \mathbb{C}[\mathfrak{h}^*]$ as vector spaces.

The Cherednik algebra $H_{t,c}(W)$ can also regarded as a deformation of the skew group ring $A_n \rtimes W$ of the Weyl algebra; in this case the spherical subalgebra $eH_{t,c}(W)e$ becomes a deformation of the fixed ring A_n^W. This is most readily seen through the Dunkl embedding of $H_c(W)$ into a localisation $\mathcal{D}(\mathfrak{h}_{\mathrm{reg}}) \rtimes W$ of $A_n \rtimes W$ (see Subsection 1.8 for the details). However, in many ways the intuition from Lie theory is more fruitful;

for example, $H_c(W)$ can have finite-dimensional representations, whereas $A_n \rtimes W$ is always simple and so cannot have any such representations.

The triangular decomposition is particularly useful for representation theory, since one has natural analogues of the Verma modules and Category $\mathcal{O}$, which are so powerful in the representation theory of semisimple Lie algebras (see, for example [83]). For Cherednik algebras, Category $\mathcal{O}$ consists of the full subcategory of finitely generated $H_c(W)$-modules on which $\mathbb{C}[\mathfrak{h}^*]$ acts locally nilpotently. The most obvious such modules are the *standard modules* $\Delta(\lambda) = H_c(W) \otimes_{\mathbb{C}[\mathfrak{h}^*]} \lambda$, where λ is an irreducible representation of W on which $\mathbb{C}[\mathfrak{h}^*]$ is given a trivial action. The structure of these modules is very similar to that of Verma modules; for example, each Δ_λ has a unique simple factor module and these define all the simple objects in Category $\mathcal{O}$. The general theory of Category $\mathcal{O}$-modules is given in Section 2. In Type A_{n-1}, when W is the symmetric group S_n, one can get a much more complete description of these modules, as is explained in Section 3. For example, it is known exactly when $H_c(S_n)$ has a finite-dimensional simple module (curiously, $H_c(S_n)$ can never have more than one such module). Moreover, the composition factors of the $\Delta(\lambda)$ and character formulae for the simple modules in Category $\mathcal{O}$ are known. The answers are given in terms of some beautiful combinatorics relating two fundamental bases of representations of certain quantum groups (more precisely, the level-one Fock spaces for quantum affine Lie algebras of Type A).

Section 4 deals with the Knizhnik–Zamolodchikov (KZ) functor. This remarkable functor allows one to relate Category $\mathcal{O}$ to modules over yet another important algebra, in this case the *cyclotomic Hecke algebra* $H_q(W)$ related to W. At its heart the KZ functor is quite easy to describe. Recall that the Dunkl embedding identifies $H_c(W)$ with a subalgebra of $\mathcal{D}(\mathfrak{h}_{\text{reg}}) \rtimes W$, and in fact $\mathcal{D}(\mathfrak{h}_{\text{reg}}) \rtimes W$ is then a localisation of $H_c(W)$. The key idea behind the KZ functor is that one can also localise the given module to obtain a $(\mathcal{D}(\mathfrak{h}_{\text{reg}}) \rtimes W)$-module. At this point powerful results from the theory of $\mathcal{D}$-modules can be applied and these results ultimately lead to modules over the Hecke algebra.

When we first defined the symplectic reflection algebras $H_{t,c}(G)$ there was the second parameter t and the representation-theoretic results we have described so far have all been concerned with the case $t \neq 0$. The case $t = 0$, which is the topic of the final Section 5, has a rather different flavour. The reason is that $H_{0,t}(G)$ is now a finite module over its centre $Z_c(G) = Z(H_{0,c}(G))$.

We again give a thorough description of the representation theory of $H_{0,t}(G)$ although this has a much more geometric flavour with a strong connection to Poisson and even symplectic geometry. The Poisson structure on Spec $Z_c(G)$ comes from the fact that the parameter t gives a quantization of $Z_c(G)$! A key observation here is that the simple $H_{0,t}(G)$-modules are finite-dimensional, of dimension bounded by $|G|$ (see Theorem 5.1.4). Moreover they have maximal dimension precisely when their central annihilator is a smooth point of Spec $Z_c(G)$. So, the geometry of that space and the representations of $H_{0,c}(G)$ are intimately connected.

The final topic in this chapter relates to Chapter IV and concerns applications of $H_{0,c}(G)$ back to algebraic geometry. An important geometric question is to understand *resolutions of singularities* $\pi : \widetilde{X} \to X$ of a singular space X (thus $\widetilde{X}$ should be nonsingular and the birational map π should be an isomorphism outside the singular subset of X). Here we are interested in the case $X = Y = V/G$. In this setting, the smooth locus of Y is even symplectic, and so one would further like the resolution $\pi : \widetilde{Y} \to Y$ to be *a symplectic resolution* in the sense that $\widetilde{Y}$ is symplectic and π is an isomorphism of symplectic spaces away from the singular set. Remarkably the question of when this happens has been determined using $H_{0,c}(G)$ — it happens if and only if Spec $Z_c(G)$ is smooth for some value of c. This can be made more precise. Recall that the spherical subalgebra $eH_{0,c}(G)e$ is a deformation of $\mathbb{C}[Y] = \mathbb{C}[V]^G$. Indeed, $eH_{0,c}(G)e \cong Z_c(G)$ is even commutative and so is a commutative deformation of $\mathbb{C}[Y]$. Thus Y has a symplectic resolution if and only if $eH_{0,c}(G)e$ is a smooth deformation of $\mathbb{C}[Y]$ (see Theorem 5.8.3). Completing this circle of ideas we note that symplectic reflection algebas have even been used to determine the groups G for which $Y = V/G$ has a symplectic resolution of singularities.

Structure of Chapter III. The aim of the chapter is to give an introduction to the construction and representation theory of symplectic reflection algebras $H_{t,c}(G)$. The basic definitions and structure theorems, including the PBW Theorem and deformation theory, are given in Section 1. In Section 2 the representation theory of the Cherednik algebra $H_c(W) = H_{1,c}(W)$ is discussed, with emphasis on Category $\mathcal{O}$-modules. In particular, Category $\mathcal{O}$ is shown to be highest weight category. These results can be considerably refined when $W = S_n$ is a symmetric group, and this case is studied in detail in Section 3. Here one can completely describe the characters of simple $\mathcal{O}$-modules and the composition factors of the standard modules. This is achieved by relating $H_c(S_n)$ to certain Schur and quantum algebras. The KZ functor is described in Section 4 and again relates the representation theory of $H_c(W)$ to other subjects: in this case the theory of $\mathcal{D}$-modules and, ultimately, to cyclotomic Hecke algebras $H_q(W)$. This allows one to prove subtle and nontrivial results about both the Cherednik and Hecke algebras. The final Section 5 studies the representation theory of the symplectic reflection algebras $H_{0,c}(G)$, with particular reference to their Poisson geometry and symplectic leaves. The application of these algebras to the theory of symplectic resolutions of quotient singularities is discussed briefly.

IV. Noncommutative resolutions. As we have just remarked, a fundamental problem in algebraic geometry is to understand the resolution of singularities $\pi : \widetilde{X} \to X$ of a singular space $X = \operatorname{Spec} R$. Even for nonsymplectic singularities, one can sometimes resolve the singularity by means of a noncommutative space and this can provide more information about the commutative resolutions. This theory is described in Chapter IV.

The resolution π is obtained by blowing up an ideal I of R related in some way to the singular subspace of Y. Unfortunately, the ideal I is not unique and even among rings R of (Krull) dimension three there are standard examples where different ideals I', I'' give rise to nonisomorphic resolutions of singularities. However, through work of

Bridgeland [42] and Bridgeland–King–Reid [44], these resolutions are closely related (more precisely, derived equivalent). The motivation for this chapter comes from work of Van den Bergh who, in [230, 231], abstracted Bridgeland's work by showing that one can find a noncommutative ring A that is actually derived equivalent to *both* these resolutions. Thus it is reasonable to think of A (or perhaps its category of modules) as a noncommutative resolution of singularities of Y. The purpose of this chapter is to explain how to construct such a ring A, to outline some of the methods that are used to extract the geometry, and to discuss the geometric applications.

For the rest of this introduction fix a commutative Gorenstein algebra R and set $X = \operatorname{Spec} R$ (in fact much of the theory works for Cohen–Macaulay rather than Gorenstein rings, but the theory is more easily explained in the Gorenstein case, and this also fits naturally with the other parts of the book). For simplicity, we assume throughout this introduction that R is also a normal, local domain. Then a *noncommutative crepant resolution* or *NCCR* for X (or R) is a ring A satisfying

(1) $A = \operatorname{End}_R(M)$ for some reflexive R-module M,
(2) A is a Cohen–Macaulay (CM) R-module, and
(3) $\operatorname{gldim} A = \dim X$.

Before discussing this definition, here is a simple but still very important example: let $G \subset \mathrm{SL}(2, \mathbb{C})$ be a finite subgroup acting naturally on $S = \mathbb{C}[x, y]$ and set $R = S^G$. Then it is easy to show that the skew group ring $A = S \rtimes G$, in the sense of Chapter III, is isomorphic to $\operatorname{End}_R(S)$, and it follows that A is a NCCR for the quotient singularity $X = \operatorname{Spec}(R) = \mathbb{C}^2/G$. This construction can be considerably generalised (to a polynomial ring in $n \geq 2$ variables, in particular), but the present case provides a rich supply of examples that we use throughout the chapter, not least because the finite subgroups of $\mathrm{SL}(2, \mathbb{C})$ are classified and easy to manipulate. See Sections 1 and 5 in particular.

Let us now explain some aspects of the definition of a NCCR; further details can be found in Section 2. First, the hypothesis that A be CM corresponds to the geometric property of crepancy (which is one reason these are called NC Crepant Resolutions). The definition of crepancy is harder to motivate, and is discussed in Section 4, but for symplectic singularities like the variety $Y = V/G$ from the last subsection, crepancy is equivalent to the resolution being symplectic. Since we want to obtain a smooth resolution of the given singular space it is very natural to require that $\operatorname{gldim} A < \infty$. Unfortunately, as also occurred in Chapter I, this is too weak an assumption in a noncommutative setting and so we require the stronger hypothesis (3). This is actually the same as demanding that all the simple A-modules have the same homological dimension and is in turn equivalent to demanding that A satisfy a nongraded version of the Artin–Schelter condition from Chapter I (see Corollary 4.6.3 for the details). So, once again, the definition is quite natural given the general philosophy of the book. Although NCCRs are not unique, they are at least Morita equivalent in dimension 2 (meaning that the categories of modules are equivalent) and derived equivalent in