

Flops and derived categories

Tom Bridgeland

Department of Mathematics and Statistics, The University of Edinburgh, King's Buildings,
Mayfield Road, Edinburgh, EH9 3JZ, UK (e-mail: tab@maths.ed.ac.uk)

Oblatum 7-IX-2000 & 15-VIII-2001
Published online: 19 November 2001 – © Springer-Verlag 2001

1. Introduction

This paper contains some applications of Fourier-Mukai techniques to problems in birational geometry. The main new idea is that flops occur naturally as moduli spaces of perverse coherent sheaves. As an application we prove

Theorem 1.1. *If X is a complex projective threefold with terminal singularities and*

$$\begin{array}{ccc} Y_1 & & Y_2 \\ & \searrow f_1 & \swarrow f_2 \\ & X & \end{array}$$

are crepant resolutions, then there is an equivalence of derived categories of coherent sheaves $D(Y_1) \longrightarrow D(Y_2)$.

The theorem implies in particular that birational Calabi-Yau threefolds have equivalent derived categories. This was conjectured to hold by A.I. Bondal and D.O. Orlov who proved some special cases in [3]. Here we prove Theorem 1.1 using the by-now standard techniques of Fourier-Mukai transforms, in particular the ideas developed in [5, 6].

For simplicity, let us suppose that Y is a non-singular, projective threefold and $f: Y \rightarrow X$ is a proper, birational morphism contracting a single rational curve $C \cong \mathbb{P}^1$ with normal bundle

$$\mathcal{N}_{C/Y} \cong \mathcal{O}_C(-1) \oplus \mathcal{O}_C(-1).$$

Using the theory of t -structures, we define an abelian category $\text{Per}(Y/X) \subset D(Y)$ whose objects we call *perverse* (or *perverse coherent*) sheaves on Y . A short exact sequence in $\text{Per}(Y/X)$ is just a triangle in $D(Y)$ whose vertices are all objects of $\text{Per}(Y/X)$. The next step is to construct moduli spaces of

perverse sheaves. To do this we introduce a stability condition. A *perverse point sheaf* is then defined to be a stable perverse sheaf which has the same numerical invariants as the structure sheaf of a point of Y .

Structure sheaves of points $y \in Y$ are objects of the category $\text{Per}(Y/X)$, and are stable for $y \in Y \setminus C$. For $y \in C$, the sheaf \mathcal{O}_y fits into the exact sequence

$$0 \longrightarrow \mathcal{O}_C(-1) \longrightarrow \mathcal{O}_C \longrightarrow \mathcal{O}_y \longrightarrow 0. \tag{1}$$

It turns out that \mathcal{O}_C is a perverse sheaf, but $\mathcal{O}_C(-1)$ is not, so that the triangle in $D(Y)$ arising from (1) does not define an exact sequence in $\text{Per}(Y/X)$. However the complex obtained by shifting $\mathcal{O}_C(-1)$ to the left by one place is a perverse sheaf, so there is an exact sequence of perverse sheaves

$$0 \longrightarrow \mathcal{O}_C \longrightarrow \mathcal{O}_y \longrightarrow \mathcal{O}_C(-1)[1] \longrightarrow 0 \tag{2}$$


which should be thought of as destabilizing \mathcal{O}_y .

Flipping the extension of perverse sheaves (2) gives stable objects of $\text{Per}(Y/X)$ fitting into an exact sequence of perverse sheaves

$$0 \longrightarrow \mathcal{O}_C(-1)[1] \longrightarrow E \longrightarrow \mathcal{O}_C \longrightarrow 0. \tag{3}$$

These perverse point sheaves E are not sheaves, indeed any such object has two nonzero homology sheaves $H_1(E) = \mathcal{O}_C(-1)$ and $H_0(E) = \mathcal{O}_C$. We shall use geometric invariant theory to construct a fine moduli space W parameterizing perverse point sheaves on X . Roughly speaking, the space W is obtained from X by replacing the rational curve C parameterising extensions (2) by another rational curve C' parameterising extensions (3).

The push-down $\mathbf{R}f_*(E)$ of a perverse point sheaf E is always the structure sheaf of a point $x \in X$, so there is a natural map $g: W \rightarrow X$. Moreover, the general point of W corresponds to the structure sheaf of a point $y \in Y \setminus C$, so g is birational. Thus there is a diagram of birational morphisms



The techniques developed in [5, 6] allow us to use the intersection theorem to show that W is non-singular, and that the universal family of perverse sheaves on $W \times Y$ induces a Fourier-Mukai transform $D(W) \rightarrow D(Y)$. An easy argument then shows that $g: W \rightarrow X$ is the flop of $f: Y \rightarrow X$. Theorem 1.1 follows from this because crepant resolutions of a terminal threefold are related by a finite chain of flops (see J. Kollár [10]).

The hard work in this paper goes into constructing the moduli space of perverse point sheaves. It turns out that the correct stability condition to impose on these objects is that they should be quotients of \mathcal{O}_Y in the category $\text{Per}(Y/X)$. Thus each perverse point sheaf E fits into an exact sequence of perverse sheaves

$$0 \longrightarrow F \longrightarrow \mathcal{O}_Y \longrightarrow E \longrightarrow 0 \tag{4}$$

and the space W is really a sort of perverse Hilbert scheme parameterising perverse quotients of \mathcal{O}_Y . The corresponding subobjects $F \subset \mathcal{O}_Y$ are simple, rank one sheaves, in general with torsion. Thus in the first place we construct a moduli space of simple sheaves F on Y and then use this space to parameterise the corresponding perverse point sheaves E .

The theory of perverse sheaves developed below is valid for any small contraction of canonical threefolds $f: Y \rightarrow X$. It seems natural to speculate that when $-K_Y$ is f -ample the resulting moduli space of perverse sheaves W is the *flip* of f . In that case one would not expect a derived equivalence between Y and W , but rather an embedding of derived categories $D(W) \hookrightarrow D(Y)$. What prevents us from proving such a result is our inability to do Fourier-Mukai on singular spaces. There is some hope that a better understanding of the mathematics surrounding the intersection theorem might allow flips to be studied in this way. This would be interesting for several reasons, not least because it would give a simpler and more conceptual proof of the existence of threefold flips. For now, however, this remains pure speculation!

The plan of the paper is as follows. Section 2 contains the basic definitions we need from the theory of triangulated categories. In Sect. 3 we define the category of perverse coherent sheaves and derive some of its basic properties. We also state Theorem 3.8 which guarantees the existence of fine moduli spaces of perverse point sheaves. In Sect. 4 we assume this result and use it to prove Theorem 1.1. The proof of Theorem 3.8 is given in Sects. 5, 6 and 7.

Acknowledgements I'd like to thank Bondal and Orlov for their wonderful hospitality in Moscow and for the many useful conversations I had with them. I particularly remember asking Bondal how Theorem 1.1 might be proved and receiving the cryptic reply 'Its all to do with t-structures' which of course turned out to be correct. The paper also benefitted from discussions I had with Alastair King and Miles Reid.

Notation. All schemes X are assumed to be of finite type over \mathbb{C} and all points are closed points. $D(X)$ denotes the unbounded derived category of coherent sheaves throughout. More precisely $D(X)$ is the subcategory of the derived category of quasi-coherent \mathcal{O}_X -modules consisting of complexes with coherent cohomology sheaves. The full subcategory of complexes with bounded cohomology sheaves is denoted $D^b(X)$. The i th cohomology sheaf of an object $E \in D(X)$ is denoted $H^i(E)$ and the i th homology sheaf by $H_i(E)$. Thus $H_i(E) = H^{-i}(E)$.

2. Admissible subcategories

This section contains some ideas from the general theory of triangulated categories. In particular, we define semi-orthogonal decompositions and t -structures. In fact, the first of these concepts is a special case of the second, but we give the definitions separately, since one tends to think of the two structures rather differently. t -structures were introduced in [1] in order to define perverse sheaves on stratified spaces. Semi-orthogonal decompositions also appear in [1] but their geometrical significance was first properly exploited by Bondal and Orlov [2,3]. We fix a triangulated category \mathcal{A} throughout, with its shift functor $T: \mathcal{A} \rightarrow \mathcal{A}: a \mapsto a[1]$.

In the context of birational geometry, the key point to note about derived categories is that performing a contraction corresponds to passing to a triangulated subcategory. More specifically one should consider so-called *admissible subcategories*.

Definition 2.1. A *right admissible subcategory* of \mathcal{A} is a full subcategory $\mathcal{B} \subset \mathcal{A}$ such that the inclusion functor $\mathcal{B} \hookrightarrow \mathcal{A}$ has a right adjoint.

Given a full subcategory $\mathcal{B} \subset \mathcal{A}$ one defines the right orthogonal $\mathcal{B}^\perp \subset \mathcal{A}$ to be the full subcategory

$$\mathcal{B}^\perp = \{a \in \mathcal{A} : \text{Hom}_{\mathcal{A}}(b, a) = 0 \text{ for all } b \in \mathcal{B}\}.$$

One can easily show [2] that if a full subcategory $\mathcal{B} \subset \mathcal{A}$ is right admissible then every object $a \in \mathcal{A}$ fits into a triangle

$$b \longrightarrow a \longrightarrow c \longrightarrow b[1],$$

with $b \in \mathcal{B}$ and $c \in \mathcal{B}^\perp$.

Definition 2.2. A *triangulated subcategory* of \mathcal{A} is a full subcategory $\mathcal{B} \subset \mathcal{A}$ which is closed under shifts, that is $\mathcal{B}[1] = \mathcal{B}$, such that any triangle $b_1 \longrightarrow b_2 \longrightarrow c \longrightarrow b_1[1]$ in \mathcal{A} with $b_1, b_2 \in \mathcal{B}$ has $c \in \mathcal{B}$ also.

Clearly the right orthogonal of a triangulated category is itself triangulated. If a triangulated subcategory $\mathcal{B} \subset \mathcal{A}$ is right admissible we say that \mathcal{A} has a *semi-orthogonal decomposition* into the subcategories $(\mathcal{B}^\perp, \mathcal{B})$; one should think of \mathcal{A} as being built up from these two smaller triangulated categories. Important examples of semi-orthogonal decompositions are given by the following result.

Proposition 2.3. *Let $f: Y \rightarrow X$ be a morphism of projective varieties such that $\mathbf{R}f_*(\mathcal{O}_Y) = \mathcal{O}_X$. Then the functor*

$$\mathbf{L}f^*: \mathbf{D}(X) \longrightarrow \mathbf{D}(Y)$$

embeds $\mathbf{D}(X)$ as a right admissible triangulated subcategory of $\mathbf{D}(Y)$.

Proof. The functor $\mathbf{L}f^*$ has the right adjoint $\mathbf{R}f_*$ and the composite $\mathbf{R}f_* \circ \mathbf{L}f^*$ is the identity on $D(X)$ by the projection formula and the assumption that $\mathbf{R}f_*(\mathcal{O}_Y) = \mathcal{O}_X$. It follows that $\mathbf{L}f^*$ is fully faithful so $D(X)$ can be identified with its image under $\mathbf{L}f^*$. Note that if X is non-singular then $\mathbf{L}f^*$ embeds $D^b(X)$ in $D^b(Y)$, but that this is no longer true when we allow singularities. \square

The Grauert-Riemenschneider vanishing theorem shows that the hypotheses of Proposition 2.3 hold whenever $f: Y \rightarrow X$ is a morphism of projective varieties such that Y has rational singularities and $-K_Y$ is f -ample.

Corollary 2.4. *Let $f: Y \rightarrow X$ be an extremal contraction of a canonical threefold. Then $D(X)$ is a right admissible triangulated subcategory of $D(Y)$.* \square

As we mentioned in the introduction, it is possible that flips also induce embeddings of derived categories. If this were true, one would be able to interpret the action of the minimal model program on a variety X as picking out some minimal admissible subcategory of $D(X)$.

Recall that an abelian category A sits inside its derived category $D(A)$ as the subcategory of complexes whose cohomology is concentrated in degree zero. There are by now plenty of examples of interesting algebraic and geometrical relationships which can be described by an equivalence of derived categories $D(A) \rightarrow D(B)$. Such equivalences will usually not arise from an equivalence of the underlying abelian categories A and B , indeed, this is why one must use derived categories. Changing perspective slightly one could think of a derived equivalence as being described by a single triangulated category with two different abelian categories sitting inside it. The theory of t -structures is the tool which allows one to see these different categories.

Definition 2.5. A t -structure on \mathcal{A} is a right admissible subcategory $\mathcal{A}^{\leq 0} \subset \mathcal{A}$ which is preserved by left shifts, that is $\mathcal{A}^{\leq 0}[1] \subset \mathcal{A}^{\leq 0}$.

Given a t -structure $\mathcal{A}^{\leq 0}$ on \mathcal{A} one defines $\mathcal{A}^{\leq i} = \mathcal{A}^{\leq 0}[-i]$ and $\mathcal{A}^{\geq i} = (\mathcal{A}^{\leq i-1})^\perp$. One also writes $\mathcal{A}^{< i} = \mathcal{A}^{\leq i-1}$ and $\mathcal{A}^{> i} = \mathcal{A}^{\geq i+1}$.

Definition 2.6. The *heart* (or *core*) of the t -structure $\mathcal{A}^{\leq 0} \subset \mathcal{A}$ is the full subcategory $\mathcal{H} = \mathcal{A}^{\leq 0} \cap \mathcal{A}^{\geq 0}$.

It was proved in [1] that the heart of a t -structure is an abelian category. Short exact sequences $0 \rightarrow a_1 \rightarrow a_2 \rightarrow a_3 \rightarrow 0$ in \mathcal{H} are determined by triangles $a_1 \rightarrow a_2 \rightarrow a_3 \rightarrow a_1[1]$ in \mathcal{A} with $a_i \in \mathcal{H}$ for all i .

The basic example is the *standard t -structure* on the derived category $D(A)$ of an abelian category A , given by

$$\begin{aligned} \mathcal{A}^{\leq 0} &= \{E \in D(A) : H^i(E) = 0 \text{ for all } i > 0\}, \\ \mathcal{A}^{\geq 0} &= \{E \in D(A) : H^i(E) = 0 \text{ for all } i < 0\}. \end{aligned}$$

The heart is the original abelian category \mathcal{A} . To give another example, suppose that $D(A) \rightarrow D(B)$ is an equivalence of derived categories. Then pulling back the standard t -structure on $D(B)$ gives a t -structure on $D(A)$ whose heart is the abelian category B .

Further examples are provided by admissible triangulated subcategories $\mathcal{B} \subset \mathcal{A}$. Any such subcategory defines a t -structure whose heart is trivial. In fact the converse is true: a t -structure $\mathcal{A}^{\leq 0} \subset \mathcal{A}$ satisfying $\mathcal{A}^{\leq 0} \cap \mathcal{A}^{\geq 0} = 0$ is actually a triangulated subcategory of \mathcal{A} . We shall not need this fact and the proof is left to the reader.

3. Perverse coherent sheaves

In this section we define the category of perverse sheaves with which we shall be working for the rest of the paper. It is objects of this category which will be naturally parameterised by the points of a flop. Let $f: Y \rightarrow X$ be a birational morphism of projective varieties. We shall make two assumptions, firstly that $\mathbf{R}f_*\mathcal{O}_Y = \mathcal{O}_X$, and secondly that f has relative dimension one. The example we have in mind is a small contraction of a canonical threefold.

Let us write $\mathcal{A} = D(Y)$ and $\mathcal{B} = D(X)$. By Proposition 2.3, we may identify \mathcal{B} with a right admissible triangulated subcategory of \mathcal{A} . Thus there is a semi-orthogonal decomposition $(\mathcal{C}, \mathcal{B})$ where

$$\mathcal{C} = \mathcal{B}^\perp = \{E \in D(Y) : \mathbf{R}f_*(E) = 0\}.$$

Note that objects of \mathcal{C} are supported on the exceptional locus of f .

Lemma 3.1. *An object $E \in D(Y)$ lies in \mathcal{C} precisely when its cohomology sheaves $H^i(E)$ lie in \mathcal{C} .*

Proof. There is a spectral sequence $\mathbf{R}^p f_* H^q(E) \implies H^{p+q} \mathbf{R}f_*(E)$ which degenerates because f has relative dimension one. \square

Note that the functor $\mathbf{R}f_*$ has the left adjoint $\mathbf{L}f^*$ and the right adjoint $f^!$. In this situation one may obtain t -structures on \mathcal{A} by glueing t -structures on \mathcal{B} and \mathcal{C} . For details see [1, 1.4.8–10] or [8, Ex. IV.4.2 (c)]. Lemma 3.1 allows one to use the standard t -structure on \mathcal{A} to induce a t -structure $\mathcal{C}^{\leq 0} = \mathcal{C} \cap \mathcal{A}^{\leq 0}$ on \mathcal{C} in the obvious way. Shifting this by an integer p and glueing it to the standard t -structure on $\mathcal{B} = D(X)$ gives a t -structure on \mathcal{A} satisfying

$$\begin{aligned} {}^p\mathcal{A}^{\leq 0} &= \{E \in \mathcal{A} : \mathbf{R}f_*(E) \in \mathcal{B}^{\leq 0} \text{ and } \text{Hom}_{\mathcal{A}}(E, C) = 0 \text{ for all } C \in \mathcal{C}^{>p}\}, \\ {}^p\mathcal{A}^{\geq 0} &= \{E \in \mathcal{A} : \mathbf{R}f_*(E) \in \mathcal{B}^{\geq 0} \text{ and } \text{Hom}_{\mathcal{A}}(C, E) = 0 \text{ for all } C \in \mathcal{C}^{<p}\}. \end{aligned}$$

The heart of this t -structure is the abelian category

$${}^p \text{Per}(Y/X) = {}^p\mathcal{A}^{\leq 0} \cap {}^p\mathcal{A}^{\geq 0}.$$

The integer p should be thought of as a choice of perversity. We shall be mainly interested in the case $p = -1$, and we refer to objects of the category

$$\text{Per}(Y/X) = {}^{-1}\text{Per}(Y/X)$$

as *perverse* (or *perverse coherent*) sheaves. The lemma below gives an explicit description of this category.

Lemma 3.2. *An object E of $D(Y)$ is a perverse sheaf if and only if the following three conditions are satisfied:*

- (a) $H_i(E) = 0$ unless $i = 0$ or 1 ,
- (b) $\mathbf{R}^1 f_* H_0(E) = 0$ and $\mathbf{R}^0 f_* H_1(E) = 0$,
- (c) $\text{Hom}_X(H_0(E), C) = 0$ for any sheaf C on Y satisfying $\mathbf{R}f_*(C) = 0$.

Proof. Suppose E is a perverse sheaf. The condition that $\mathbf{R}f_*(E)$ is a sheaf on X , together with the spectral sequence of Lemma 3.1, gives condition (b) and implies that $\mathbf{R}f_* H_i(E) = 0$ unless $i = 0$ or 1 .

Let $\tau_{<i}$ and $\tau_{>i}$ be the truncation functors of the standard t -structure on $D(Y)$. There are natural maps $\tau_{<i} E \rightarrow E$ and $E \rightarrow \tau_{>i} E$. Then $\tau_{>0} E = 0$ because $\tau_{>0} E \in \mathcal{C}^{>0}$. Similarly $\tau_{<-1} E = 0$ because $\tau_{<-1} E \in \mathcal{C}^{<-1}$. This proves condition (a). Condition (c) is clear, since any nonzero map from $H_0(E)$ to a sheaf $C \in \mathcal{C}^{>-1}$ induces a nonzero morphism $E \rightarrow C$ in $D(Y)$.

The converse is easy and is left to the reader. □

Definition 3.3. We shall say that two objects A_1 and A_2 of $D^b(Y)$ are *numerically equivalent* if for any locally-free sheaf L on Y one has $\chi(L, A_1) = \chi(L, A_2)$.

Recall that for objects L and A of $D^b(Y)$ with L of finite homological dimension

$$\chi(L, A) = \sum_i (-1)^i \dim \text{Hom}_{D(Y)}^i(L, A).$$

Thus if Y is a non-singular projective variety, then by the Riemann-Roch theorem, two objects of $D^b(Y)$ are numerically equivalent precisely when they have the same Chern character.

Definition 3.4. An object F of $D(Y)$ is a *perverse ideal sheaf* if there is an injection $F \hookrightarrow \mathcal{O}_Y$ in the category $\text{Per}(Y/X)$. An object E of $D(Y)$ is a *perverse structure sheaf* if there is a surjection $\mathcal{O}_Y \twoheadrightarrow E$ in the category $\text{Per}(Y/X)$. A *perverse point sheaf* is a perverse structure sheaf which is numerically equivalent to the structure sheaf of a point $y \in Y$.

Thus a perverse ideal sheaf F determines and is determined by a perverse structure sheaf E , which fit together in an exact sequence of perverse sheaves

$$0 \longrightarrow F \longrightarrow \mathcal{O}_Y \longrightarrow E \longrightarrow 0. \tag{5}$$

Applying the cohomology functor to the above exact sequence shows that perverse ideal sheaves are actually sheaves, that is satisfy $H_i(F) = 0$ for $i \neq 0$.

Example 3.5. Let us suppose, as in the introduction, that $f: Y \rightarrow X$ is the contraction of a non-singular rational curve C with normal bundle $\mathcal{O}_C(-1) \oplus \mathcal{O}_C(-1)$ on a non-singular projective threefold Y . For any point $y \in C$ there is an exact sequence of sheaves

$$0 \longrightarrow \mathcal{I}_C \longrightarrow \mathcal{I}_y \longrightarrow \mathcal{O}_C(-1) \longrightarrow 0.$$

This shows that \mathcal{I}_y is not a perverse sheaf, and it follows that although \mathcal{O}_y is a perverse sheaf, it is not a quotient of \mathcal{O}_Y in $\text{Per}(Y/X)$. Consider instead non-trivial extensions of the form

$$0 \longrightarrow \mathcal{O}_C(-1) \longrightarrow F \longrightarrow \mathcal{I}_C \longrightarrow 0.$$

One can easily calculate that $\text{Ext}_Y^1(\mathcal{I}_C, \mathcal{O}_C(-1)) = \mathbb{C}^2$, so the set of such sheaves F is parameterised by a rational curve. Composing the map $F \rightarrow \mathcal{I}_C$ with the inclusion $\mathcal{I}_C \subset \mathcal{O}_Y$ gives a nonzero morphism $F \rightarrow \mathcal{O}_Y$ and we take E to be its cone. In this way we obtain an exact sequence of perverse sheaves

$$0 \longrightarrow F \longrightarrow \mathcal{O}_Y \longrightarrow E \longrightarrow 0$$

with $H_1(E) = \mathcal{O}_C(-1)$ and $H_0(E) = \mathcal{O}_C$. Thus E is a perverse point sheaf.

The flop of Y along C is a non-singular threefold W with a morphism $g: W \rightarrow X$ contracting a single rational curve C' . We shall show that the points of W parameterise perverse point sheaves on Y . The perverse point corresponding to a point $w \in W \setminus C'$ is a point $y \in Y \setminus C$, whereas the points of C' correspond to the perverse point sheaves E described above.

Lemma 3.6. *Let E_1 and E_2 be perverse point sheaves on Y . Then*

$$\text{Hom}_{\text{D}(Y)}(E_1, E_2) = \begin{cases} \mathbb{C} & \text{if } E_1 = E_2, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. If E is a perverse point sheaf then $\mathbf{R}f_*(E)$ is the structure sheaf of a point of X , so $\text{Hom}_{\text{D}(Y)}(\mathcal{O}_Y, E) = \mathbb{C}$. Taking Homs of the exact sequence (5) into E shows that $\text{Hom}_{\text{D}(Y)}(E, E) = \mathbb{C}$.

Suppose there is a nonzero morphism $\theta: E_1 \rightarrow E_2$. Taking Homs of (5) into E_2 shows that the unique map $\mathcal{O}_Y \rightarrow E_2$ must factor via θ . In particular, θ is surjective in $\text{Per}(Y/X)$. But then the kernel K of θ in $\text{Per}(Y/X)$ is numerically equivalent to zero, and this implies $K = 0$, so θ is an isomorphism. \square

Let S be a scheme. Given a point $s \in S$, let $j_s: \{s\} \times Y \hookrightarrow S \times Y$ be the embedding. A family of sheaves on Y over S is just an object \mathcal{F} of $\text{D}(S \times Y)$ such that for each point $s \in S$ the object $\mathcal{F}_s = \mathbf{L}j_s^*(\mathcal{F})$ of $\text{D}(Y)$ is a sheaf. Indeed, by [4, Lemma 4.3], this condition implies that \mathcal{F} is actually a sheaf on $S \times Y$, flat over S , so that \mathcal{F} defines a family of sheaves in the usual sense. Once this observation has been made it is clear what the correct definition of a family of perverse sheaves should be.

Definition 3.7. A family of perverse sheaves on Y over a scheme S is an object \mathcal{E} of $D(S \times Y)$ such that for each point $s \in S$ the object $\mathcal{E}_s = \mathbf{L}j_s^*(\mathcal{E})$ of $D(Y)$ is a perverse sheaf. Two such families \mathcal{E}_1 and \mathcal{E}_2 are equivalent if $\mathcal{E}_2 = \mathcal{E}_1 \otimes L$ for some line bundle L pulled back from S .

The proof of the following theorem will be given in Sects. 5 and 6 below.

Theorem 3.8. *The functor which assigns to a scheme S the set of equivalence classes of families of perverse point sheaves on Y over S is representable by a projective scheme $\mathcal{M}(Y/X)$.*

We conclude this section with the following base-change result.

Proposition 3.9. *Let S be a scheme and \mathcal{E} a family of perverse sheaves on Y over S . Put $f_S = \text{id}_S \times f$. Then $\mathcal{G} = f_{S*}(\mathcal{E})$ is a family of sheaves on X over S , and for any point $s \in S$ there is an isomorphism of sheaves $\mathcal{G}_s = \mathbf{R}f_*(\mathcal{E}_s)$.*

Proof. Fix a point $s \in S$ and consider the diagram

$$\begin{array}{ccc} Y & \xhookrightarrow{j_s} & S \times Y \\ \downarrow f & & \downarrow f_S \\ X & \xhookrightarrow{i_s} & S \times X \end{array}$$

where $i_s: \{s\} \times X \hookrightarrow S \times X$ is the embedding.

If $p: S \times X \rightarrow S$ is the projection map, flat base-change shows that $i_{s*}\mathcal{O}_X = p^*\mathcal{O}_S$. It follows that $\mathbf{L}f_S^* \circ i_{s*}\mathcal{O}_X = j_{s*}\mathcal{O}_Y$. The projection formula gives isomorphisms

$$\begin{aligned} i_{s*} \circ \mathbf{R}f_*(\mathcal{E}_s) &= \mathbf{R}f_{S*} \circ j_{s*} \circ \mathbf{L}j_s^*(\mathcal{E}) = \mathbf{R}f_{S*}(j_{s*}\mathcal{O}_Y \overset{\mathbf{L}}{\otimes} \mathcal{E}) \\ &= i_{s*}(\mathcal{O}_X) \overset{\mathbf{L}}{\otimes} \mathbf{R}f_{S*}(\mathcal{E}) = i_{s*} \circ \mathbf{L}i_s^*(\mathcal{G}). \end{aligned}$$

The functor i_{s*} is exact and fully faithful on the category of sheaves on X . Thus $\mathbf{R}f_*(\mathcal{E}_s)$ is a sheaf precisely when $\mathbf{L}i_s^*(\mathcal{G})$ is. Since \mathcal{E} is a family of perverse sheaves, $\mathbf{R}f_*(\mathcal{E}_s)$ is a sheaf for all $s \in S$, so \mathcal{G} is a sheaf, flat over S . The result follows. □

4. Flops and the derived equivalence

In this section we shall show how fine moduli spaces of perverse point sheaves give rise to Fourier-Mukai type equivalences of derived categories. To do this we shall assume that fine moduli spaces of perverse point sheaves exist (as in Theorem 3.8) and apply the techniques of [5,6]. In this way we obtain a proof of Theorem 1.1.

(4.1) Let X be a projective threefold with Gorenstein, terminal singularities. Recall that a *crepant resolution* is a morphism $f: Y \rightarrow X$ from a non-singular projective variety Y , such that $f^*\omega_X = \omega_Y$. Any such resolution satisfies $\mathbf{R}f_*(\mathcal{O}_Y) = \mathcal{O}_X$ and contracts only a finite number of curves. Thus the open subset $U \subset X$ over which f is an isomorphism is the complement of a finite set of points.


By Theorem 3.8 there is a fine moduli space $\mathcal{M}(Y/X)$ of perverse point sheaves on Y . Each point $y \in f^{-1}(U)$ is a perverse point sheaf so there is an embedding $U \hookrightarrow \mathcal{M}(Y/X)$. Let $W \subset \mathcal{M}(Y/X)$ be the irreducible component of $\mathcal{M}(Y/X)$ containing the image of this morphism. In fact, it is possible to prove, as in [5, Sect. 8], that $\mathcal{M}(Y/X)$ is irreducible, so that $W = \mathcal{M}(Y/X)$, but we shall not need this.

Let \mathcal{P} be a universal object on $W \times Y$. Thus \mathcal{P} is an object of $D(W \times Y)$ such that the perverse point sheaf on Y corresponding to a point $w \in W$ is the object $\mathcal{P}_w = \mathbf{L}i_w^*(\mathcal{P})$, where $i_w: \{w\} \times Y \hookrightarrow W \times Y$ is the embedding.

(4.2) By Proposition 3.9, the sheaf $\mathbf{R}(\text{id}_W \times f)_*(\mathcal{P})$ is a family of structure sheaves of points on X over W , and therefore, up to a twist by the pullback of a line bundle from W , is the structure sheaf of the graph $\Gamma(g) \subset W \times X$ of some morphism $g: W \rightarrow X$. Thus twisting \mathcal{P} by the pullback of a line bundle from W , we can assume that

$$\mathbf{R}(\text{id}_W \times f)_*(\mathcal{P}) = \mathcal{O}_{\Gamma(g)}. \tag{6}$$

With this condition \mathcal{P} is uniquely defined. The morphism g is birational because for any point $x \in U$ there is only one object E of $D(Y)$ satisfying $\mathbf{R}f_*(E) = \mathcal{O}_x$. Thus there is a diagram of birational morphisms



(4.3) The scheme W is a non-singular projective variety and $g: W \rightarrow X$ is a crepant resolution. Furthermore, the Fourier-Mukai functor

$$\Phi(-) = \mathbf{R}\pi_{Y,*}(\mathcal{P} \otimes^{\mathbf{L}} \pi_W^*(-)): D(W) \longrightarrow D(Y),$$

is an equivalence of categories which takes $D^b(W)$ into $D^b(Y)$.

Proof. Each object \mathcal{P}_w has bounded homology sheaves, and Y is non-singular, so the object \mathcal{P} has finite homological dimension. It follows that the functor Φ takes $D^b(W)$ into $D^b(Y)$.

For each point $w \in W$ the object \mathcal{P}_w is simple, so its support is connected, and since $\mathbf{R}f_*(\mathcal{P}_w) = \mathcal{O}_x$, where $x = g(w)$, it follows that \mathcal{P}_w is supported on the fibre of f over x . Since f is crepant this implies that $\mathcal{P}_w \otimes_{\mathcal{O}_Y} = \mathcal{P}_w$.

Given distinct points $w_1, w_2 \in W$, Serre duality together with Lemma 3.6 shows that

$$\text{Hom}_{\mathbf{D}(Y)}^i(\mathcal{P}_{w_1}, \mathcal{P}_{w_2}) = 0$$

unless $g(w_1) = g(w_2)$ and $1 \leq i \leq 2$. The argument of [5, Sect. 6] then implies that W is non-singular, g is crepant and Φ is an equivalence. \square

(4.4) An immediate consequence of the isomorphism (6) is that there is a commutative diagram of functors

$$\begin{array}{ccc} \mathbf{D}(W) & \xrightarrow{\Phi} & \mathbf{D}(Y) \\ & \searrow \mathbf{R}g_* & \swarrow \mathbf{R}f_* \\ & & \mathbf{D}(X) \end{array}$$

(4.5) If C is a sheaf on W satisfying $\mathbf{R}g_*(C) = 0$ then $\Phi(C)[-1]$ is a sheaf on Y .

Proof. First suppose that A is an object of $\mathbf{D}(W)$ with $\mathbf{R}g_*(A) = 0$ and such that $B = \Phi(A)[-1]$ is a sheaf on Y . By (4.4) one has $\mathbf{R}f_*(B) = 0$ so by Lemma 3.2 the object $B[1]$ is a perverse sheaf on Y . Thus for any point $w \in W$,

$$\text{Hom}_{\mathbf{D}(W)}^i(A, \mathcal{O}_w) = \text{Hom}_{\mathbf{D}(Y)}^i(B[1], \mathcal{P}_w) = 0 \text{ unless } 0 \leq i \leq 3.$$

Moreover, since \mathcal{P}_w is a quotient of \mathcal{O}_Y in $\text{Per}(Y/X)$,

$$\text{Hom}_{\mathbf{D}(Y)}^3(B[1], \mathcal{P}_w) = \text{Hom}_{\mathbf{D}(Y)}^0(\mathcal{P}_w, B[1]) = 0.$$

Thus the object A has homological dimension at most two, and is supported in codimension at least two. It follows from this that A is a sheaf on W (see, for example, [6, Lemma 4.2]).

Now assume that C is a sheaf on W satisfying $\mathbf{R}g_*(C) = 0$ and suppose that $D = \Phi(C)[-1]$ is not a sheaf on Y . As in the proof of Lemma 3.2, we can find a sheaf B on Y satisfying $\mathbf{R}f_*(B) = 0$, and an integer $i < 0$, such that one of $\text{Hom}_{\mathbf{D}(Y)}^i(B, D)$ or $\text{Hom}_{\mathbf{D}(Y)}^i(D, B)$ is nonzero. Since Φ is an equivalence, $B = \Phi(A)[-1]$ for some object A of $\mathbf{D}(W)$, and by the first part A is a sheaf. This implies that one of the spaces $\text{Hom}_{\mathbf{D}(W)}^i(A, C)$ or $\text{Hom}_{\mathbf{D}(W)}^i(C, A)$ is nonzero, which is impossible since A and C are both sheaves. \square

(4.6) The variety $W = \mathcal{M}(Y/X)$ is the flop of $f: Y \rightarrow X$, that is, if D is a divisor on W such that $-D$ is g -nef, then its proper transform D' on Y is f -nef.

Proof. Let C be a rational curve on W contracted by g . Put $M = \Phi(\mathcal{O}_W(D))$ and $N = \Phi(\mathcal{O}_C(-1))$. Then by Riemann-Roch,

$$\chi(M, N) = \chi(\mathcal{O}_W(D), \mathcal{O}_C(-1)) = -D \cdot C \geq 0.$$

Over the open subset $f^{-1}(U)$ of Y , M is isomorphic to $\mathcal{O}_Y(D')$, and it follows that $c_1(M) = [D']$. By (4.5) the object $N[-1]$ is a sheaf supported on some curve C' of Y which is contracted by f . It follows that $\chi(M, N) = D' \cdot C'$ and hence the result. □

(4.7) Define full triangulated subcategories

$$\mathcal{C}(W/X) \hookrightarrow D(W) \quad \mathcal{C}(Y/X) \hookrightarrow D(Y)$$

consisting of objects satisfying $\mathbf{R}g_*(C) = 0$ and $\mathbf{R}f_*(C) = 0$ respectively. These categories inherit t -structures from the standard t -structures on $D(W)$ and $D(Y)$, as in Sect. 3. There is a commutative diagram of functors

$$\begin{array}{ccccc} \mathcal{C}(W/X) & \hookrightarrow & D(W) & \xrightarrow{\mathbf{R}g_*} & D(X) \\ \downarrow \tilde{\Phi} & & \downarrow \Phi & & \downarrow \text{id} \\ \mathcal{C}(Y/X) & \hookrightarrow & D(Y) & \xrightarrow{\mathbf{R}f_*} & D(X) \end{array}$$

in which the rows are exact sequences of triangulated categories. By (4.5), the equivalence $\tilde{\Phi}[-1]$ is t -exact, that is preserves the t -structures.

(4.8) There is a chain of exact equivalences of abelian categories

$$\begin{array}{ccccccc} \dots & \longrightarrow & {}^{-1}\text{Per}(W/X) & \longrightarrow & {}^0\text{Per}(Y/X) & \longrightarrow & {}^1\text{Per}(W/X) \\ & & & & & & \\ & & & & & & \longrightarrow {}^2\text{Per}(Y/X) \longrightarrow \dots \end{array}$$

Indeed, it follows from (4.7) that for any integer p the functor Φ induces an exact equivalence

$${}^p\text{Per}(W/X) \cong {}^{p+1}\text{Per}(Y/X),$$

and since the flopping operation is an involution we may interchange Y and W .

5. Perverse ideal sheaves

In this section we use geometric invariant theory to construct fine moduli spaces of perverse ideal sheaves. As in Sect. 3, let $f: Y \rightarrow X$ be a birational morphism of projective varieties of relative dimension one and satisfying $\mathbf{R}f_*(\mathcal{O}_Y) = \mathcal{O}_X$. Our first task is to identify which objects of $D(Y)$ are perverse ideal sheaves.

Proposition 5.1. *A perverse ideal sheaf on Y is, in particular, a sheaf on Y . Furthermore, a sheaf F on Y is a perverse ideal sheaf if and only if the following two conditions are satisfied:*

- (a) *the sheaf $f_*(F)$ on X is an ideal sheaf,*
- (b) *the natural map of sheaves $\eta: f^*f_*(F) \rightarrow F$ is surjective.*

Proof. Let F be a perverse ideal sheaf on Y and E the corresponding perverse structure sheaf. Applying the homology functor to the exact sequence (5) and using Lemma 3.2 shows that F is a sheaf. The functor $\mathbf{R}f_*$ is exact on the category of perverse sheaves, so there is an exact sequence of sheaves

$$0 \longrightarrow f_*(F) \longrightarrow \mathcal{O}_X \longrightarrow \mathbf{R}f_*(E) \longrightarrow 0.$$

It follows that $f_*(F)$ is an ideal sheaf on X .

Let A, B and C denote the kernel, cokernel and image of the map η in the category of sheaves on Y . Thus we have a pair of short exact sequences fitting into a diagram

$$\begin{array}{ccccccc}
 0 & \longrightarrow & A & \longrightarrow & f^*f_*F & \xrightarrow{\eta} & F & \longrightarrow & B & \longrightarrow & 0. \\
 & & & & & \searrow & \nearrow & & & & \\
 & & & & & & C & & & &
 \end{array}$$

The spectral sequence of Lemma 3.1 gives an exact sequence

$$0 \longrightarrow \mathbf{R}^1f_*(\mathbf{L}_1f^*f_*(F)) \longrightarrow f_*(F) \longrightarrow f_*(f^*f_*(F)) \longrightarrow 0,$$

together with the fact that $\mathbf{R}^1f_*(f^*f_*(F)) = 0$. Since $f_*(F)$ is torsion-free this implies that $f_*(f^*f_*(F)) = f_*(F)$.

The morphism f is birational, so η is generically an isomorphism and A and B are torsion sheaves. Applying f_* to the exact sequences above shows that $f_*(F)$ injects into $f_*(C)$ and also $f_*(C)$ injects into $f_*(F)$. It follows that $f_*(C) = f_*(F)$ and so $f_*(B) = 0$. Since F is perverse, $\mathbf{R}^1f_*(F) = 0$, so $\mathbf{R}f_*(B) = 0$ and hence by Lemma 3.2, $B = 0$, that is, η is surjective.

For the converse, suppose F is a sheaf on Y satisfying our two conditions. There is an exact sequence of sheaves

$$0 \longrightarrow A \longrightarrow f^*f_*(F) \xrightarrow{\eta} F \longrightarrow 0. \tag{7}$$

It follows that $\mathbf{R}^1f_*(F) = 0$ and $\text{Hom}_Y(F, C) = 0$ for any sheaf C on Y satisfying $\mathbf{R}f_*(C) = 0$, hence, by Lemma 3.2, F is a perverse sheaf on Y . Since f is birational, η is generically an isomorphism, so A is a torsion sheaf and

$$\text{Hom}_Y(F, \mathcal{O}_Y) = \text{Hom}_Y(f^*f_*(F), \mathcal{O}_Y) = \text{Hom}_X(f_*(F), \mathcal{O}_X)$$

which is nonzero because $f_*(F)$ is an ideal sheaf. Take a nonzero morphism $F \rightarrow \mathcal{O}_Y$ and form a triangle

$$F \longrightarrow \mathcal{O}_Y \longrightarrow E \longrightarrow F[1].$$

It will be enough to show that $E \in \text{Per}(Y/X)$. Applying the homology functor gives a long exact sequence of sheaves

$$0 \longrightarrow H_1(E) \longrightarrow F \longrightarrow \mathcal{O}_Y \longrightarrow H_0(E) \longrightarrow 0.$$

It follows that $\mathbf{R}^1 f_*(H_0(E)) = 0$ and $\text{Hom}_Y(H_0(E), C) = 0$ for any sheaf C on Y satisfying $\mathbf{R}f_*(C) = 0$. Furthermore $f_*(F)$ is torsion-free so $\mathbf{R}^0 f_*(H_1(E)) = 0$. Applying Lemma 3.2 completes the proof. \square

Lemma 5.2. *Any perverse ideal sheaf F on Y is simple, that is $\text{Hom}_Y(F, F) = \mathbb{C}$.*

Proof. The ideal sheaf $f_*(F)$ is simple, so $\text{Hom}_Y(f^*f_*(F), F) = \mathbb{C}$. Applying the functor $\text{Hom}_Y(-, F)$ to the sequence (7) gives the result. \square

6. Moduli of perverse sheaves

In this section we construct a fine moduli space of perverse ideal sheaves. To do this we mimic C. Simpson’s proof [12, Sect. 1] of the existence of moduli spaces of semistable sheaves on projective schemes. As before $f : Y \rightarrow X$ denotes a birational morphism of projective varieties of relative dimension one satisfying $\mathbf{R}f_*(\mathcal{O}_Y) = \mathcal{O}_X$.

A very special case of Simpson’s construction shows that if one takes a sufficiently ample line bundle $\mathcal{O}_Y(1)$ on Y and a suitable vector space V , then the stable points for the action of the group $\text{SL}(V)$ acting on the Quot scheme parameterising quotients of $V \otimes \mathcal{O}_Y(-1)$ with rank one and trivial determinant are just the points corresponding to ideal sheaves on Y . Here we shall show that if one takes a sufficiently ample line bundle $\mathcal{O}_X(1)$ on X and replaces $\mathcal{O}_Y(-1)$ by $f^*\mathcal{O}_X(-1)$ in the above construction, then the stable points are precisely the points corresponding to perverse ideal sheaves on Y .

Throughout we shall fix a numerical equivalence class (γ) . Let F denote a perverse ideal sheaf on Y in this class.

Rank one, torsion-free sheaves in a given numerical equivalence class form a bounded family, so we may choose $\mathcal{O}_X(1)$ so that for any torsion-free sheaf A on X in the same numerical equivalence class as $f_*(F)$, the sheaf $A \otimes \mathcal{O}_X(1)$ is generated by its global sections and satisfies $H^i(X, A \otimes \mathcal{O}_X(1)) = 0$ for all $i > 0$.

Put $L = f^*(\mathcal{O}_X(-1))$ and let V be the vector space

$$V = \text{Hom}_Y(L, F) = \text{Hom}_X(\mathcal{O}_X(-1), f_*(F)).$$

Let Quot denote the Quot scheme parameterising quotients of the vector bundle $V \otimes L$ in the numerical equivalence class (γ) .

Fix a very ample line bundle $\mathcal{O}_Y(1)$ on Y . For sufficiently large integers m , there is a closed embedding of Quot in a Grassmannian, which sends the quotient $V \otimes L \rightarrow F$ to the quotient of vector spaces

$$\text{Hom}_Y(\mathcal{O}_Y(-m), V \otimes L) \rightarrow \text{Hom}_Y(\mathcal{O}_Y(-m), F).$$

This embedding is $\mathrm{SL}(V)$ -equivariant. By pulling back the natural $\mathrm{SL}(V)$ -equivariant polarisation of the Grassmannian one obtains an $\mathrm{SL}(V)$ -equivariant polarisation $\mathcal{L}(m)$ of Quot . The following two results of Simpson describe the stable points of this group action.

Lemma 6.1. [12, Lemma 1.15] *There exists M such that for $m \geq M$ the following holds. Suppose $V \otimes L \twoheadrightarrow F$ is a point in Quot . For any subspace $H \subset V$, let G denote the subsheaf of F generated by $H \otimes L$. Suppose that $P_G(m) > 0$ and*

$$\frac{\dim(H)}{P_G(m)} < \frac{\dim(V)}{P_F(m)}$$

for all nonzero proper subspaces H . Then the point is stable with respect to the polarisation $\mathcal{L}(m)$ and the action of $\mathrm{SL}(V)$. \square

Lemma 6.2. [12, Lemma 1.16] *There exists M such that for $m \geq M$, if F is the quotient sheaf represented by a point of Quot which is semistable with respect to \mathcal{L}_m and the action of $\mathrm{SL}(V)$, then for any nonzero subspace $H \subset V$ the subsheaf of F generated by $H \otimes L$ has positive rank. \square*

Next we pick out that part of Quot which parameterises quotients which are perverse ideal sheaves.

Lemma 6.3. *There is a locally-closed subscheme $\mathcal{U} \subset \mathrm{Quot}$ parameterising quotients $V \otimes L \twoheadrightarrow F$ for which $f_*(F)$ is an ideal sheaf and for which the natural map $\alpha_F: V \rightarrow \mathrm{Hom}_Y(L, F)$ is an isomorphism.*

Proof. By Proposition 3.9 there is an open subscheme of Quot parameterising quotients for which $f_*(F)$ is torsion-free. There is also a closed subscheme over which there is a nonzero map $f_*(F) \rightarrow \mathcal{O}_X$. These two conditions are equivalent to $f_*(F)$ being an ideal sheaf. The condition on the map α_F is clearly open. \square

Let Quot^0 be the closure of \mathcal{U} in Quot . The action of $\mathrm{SL}(V)$ on Quot preserves \mathcal{U} and hence Quot^0 so we obtain an induced equivariant polarisation of Quot^0 which we also denote by $\mathcal{L}(m)$. As noted by Simpson [12, Proposition 1.12] the stable points for the induced action of $\mathrm{SL}(V)$ with respect to $\mathcal{L}(m)$ are just those stable points of the action of $\mathrm{SL}(V)$ on Quot with respect to $\mathcal{L}(m)$ which lie in Quot^0 .

Proposition 6.4. *For all sufficiently large integers m the following result holds. The semistable points for the action of $\mathrm{SL}(V)$ on the projective scheme Quot^0 with respect to the polarisation $\mathcal{L}(m)$ are precisely the points of the open subset \mathcal{U} . Furthermore all these points are properly stable.*

Proof. Let m be larger than the integer M given by Lemma 6.1 and take a point $p \in \mathcal{U}$ corresponding to a quotient $V \otimes L \twoheadrightarrow F$. Let $H \subset V$ be a proper nonzero subspace and let $G \subset F$ be the subsheaf generated by $H \otimes L$. Since α_F is an isomorphism, G is nonzero. But G is a quotient of $H \otimes L$, so $f_*(G)$ is nonzero, and since $f_*(F)$ is torsion-free this implies that G has rank one. Thus the leading coefficient of the Hilbert polynomial P_G is the same as that of P_F .

Since the class of subsheaves of F which are quotients of $V \otimes L$ is bounded, the set of possible Hilbert polynomials of G is finite, so we may assume that the inequality

$$\frac{\dim(H)}{P_G(m)} < \frac{\dim(V)}{P_F(m)}$$

holds. Lemma 6.1 then implies that the point p is stable with respect to $\mathcal{L}(m)$.

For the converse, let m be larger than the integer M given by Lemma 6.2 and take a point $p \in \text{Quot}^0$ which is semistable with respect to $\mathcal{L}(m)$. Let $V \otimes L \twoheadrightarrow F$ be the corresponding quotient. By Lemma 6.2, for any nonzero subspace $H \subset V$, the subsheaf generated by $H \otimes L$ has positive rank. It follows that the map α_F is an isomorphism.

Let T be the torsion subsheaf of $f_*(F)$ and put $Q = f_*(F)/T$. Let $H \subset V$ be the subspace $\text{Hom}_X(\mathcal{O}_X(-1), T)$. If H is nonzero $H \otimes L$ generates a rank one subsheaf of F . Since f is birational this contradicts the fact that T is torsion. Thus $H = 0$ and the natural map $V \rightarrow \text{Hom}_X(\mathcal{O}_X(-1), Q)$ is injective.

By assumption $f_*(F)$ is a degeneration of torsion-free sheaves, so by [12, Lemma 1.17] there exists a torsion-free sheaf A with the same numerical invariants as $f_*(F)$ and an inclusion $Q \subset A$. By our choice of $\mathcal{O}_X(1)$ the vector space $\text{Hom}_X(\mathcal{O}_X(-1), A)$ has the same dimension as V and the sheaf $A \otimes \mathcal{O}_X(1)$ is generated by its global sections. It follows that $Q = A$. Then T is numerically trivial, so $T = 0$ and $f_*(F)$ is torsion-free. By semi-continuity there is a nonzero map $f_*(F) \rightarrow \mathcal{O}_X$, so $f_*(F)$ is an ideal sheaf. □

Theorem 6.5. *The functor which assigns to a scheme S the set of equivalence classes of families over S of perverse ideal sheaves in a given numerical equivalence class (γ) is representable by a projective scheme $\mathcal{M}_{PI}(Y/X; \gamma)$.*

Proof. Let F be a perverse ideal sheaf on Y in the numerical equivalence class (γ) . By assumption $f_*(F) \otimes \mathcal{O}_X(1)$ is generated by its global sections so there exists a surjection $V \otimes \mathcal{O}_X(-1) \twoheadrightarrow f_*(F)$. Pulling back and using Proposition 5.1 shows that there is a surjection $V \otimes L \twoheadrightarrow F$ and hence a point of \mathcal{U} for which the corresponding quotient is F .

Conversely, the argument of Lemma 3.2 shows that a quotient $V \otimes L \twoheadrightarrow F$ corresponding to a point of \mathcal{U} is a perverse ideal sheaf. Exactly as

in [12, Theorem 1.21] we can conclude that there is a coarse moduli space $\mathcal{M}_{PI}(Y/X; \gamma)$ for perverse ideal sheaves in the numerical equivalence class (γ) , and that a universal sheaf exists locally in the étale topology on \mathcal{M} .

All perverse ideal sheaves F have rank one, so $\chi(F, \mathcal{O}_y) = 1$ for any point $y \in Y$. If y is a non-singular point then \mathcal{O}_y has a finite locally-free resolution, so the integers $\chi(F \otimes L)$, as L ranges over all locally-free sheaves on Y , have no common factor. Since the sheaves F are simple, an argument of S. Mukai [11, Theorem A.6] shows that one can patch the local universal sheaves to obtain a universal sheaf on $\mathcal{M}_{PI}(Y/X; \gamma)$. This completes the proof. \square

7. Perverse Hilbert schemes

In this section we complete the proof of Theorem 3.8. To do this we construct a perverse Hilbert scheme $\text{P-Hilb}(Y/X)$ parameterising quotients of \mathcal{O}_Y in the category $\text{Per}(Y/X)$. Again $f: Y \rightarrow X$ denotes a birational morphism of projective varieties of relative dimension one and satisfying $\mathbf{R}f_*(\mathcal{O}_Y) = \mathcal{O}_X$.

If S is a scheme, the S -valued points of the Hilbert scheme of a variety Y consist of triangles $\mathcal{F} \rightarrow \mathcal{O}_{S \times Y} \rightarrow \mathcal{E} \rightarrow \mathcal{F}[1]$ in $\text{D}(S \times Y)$ such that \mathcal{E} and \mathcal{F} are families of sheaves on Y over S . Analogously we make the

Definition 7.1. Let $\text{P-Hilb}(Y/X)$ be the functor which assigns to a scheme S the set of isomorphism classes of triangles $\mathcal{F} \rightarrow \mathcal{O}_{S \times Y} \rightarrow \mathcal{E} \rightarrow \mathcal{F}[1]$ in $\text{D}(S \times Y)$ such that \mathcal{F} and \mathcal{E} define families of perverse sheaves on Y over S .

Given a scheme S we write $f_S = \text{id}_S \times f: S \times Y \rightarrow S \times X$.

Lemma 7.2. *If \mathcal{F} is a family of perverse ideal sheaves on Y over S , the map*

$$f_{S*}: \text{Hom}_{S \times Y}(\mathcal{F}, \mathcal{O}_{S \times Y}) \rightarrow \text{Hom}_{S \times X}(f_{S*}(\mathcal{F}), \mathcal{O}_{S \times X})$$

is an isomorphism.

Proof. Since perverse ideal sheaves are sheaves, the object \mathcal{F} is a sheaf on $S \times Y$, flat over S . Consider the natural map of sheaves $\eta: f_S^* f_{S*}(\mathcal{F}) \rightarrow \mathcal{F}$. For each point $s \in S$, the map $\eta_s = \mathbf{L}j_s^*(\eta)$ is just the natural map $f^* f_*(\mathcal{F}_s) \rightarrow \mathcal{F}_s$ which is surjective by Proposition 5.1. It follows that η is surjective.

Let K be the kernel of η . It will be enough to show that there are no nonzero maps $K \rightarrow \mathcal{O}_{S \times Y}$. For this we may assume that S is affine. Let $p: S \times Y \rightarrow Y$ be the projection. Since f is birational, $p_*(K)$ is a torsion sheaf. But by base-change $p_*(\mathcal{O}_{S \times Y}) = H^0(S, \mathcal{O}_S) \otimes \mathcal{O}_Y$ is torsion-free, so any map $p_*(K) \rightarrow p_*(\mathcal{O}_{S \times Y})$ is zero. \square

The functor $\text{P-Hilb}(Y/X)$ decomposes into components which parameterise ideal sheaves in a given numerical equivalence class (γ) :

$$\text{P-Hilb}(Y/X) = \coprod_{(\gamma)} \text{P-Hilb}(Y/X; \gamma).$$

Theorem 7.3. *For each numerical equivalence class (γ) the corresponding functor $\text{P-Hilb}(Y/X; \gamma)$ is representable by a projective scheme.*

Proof. By Theorem 6.5 there is a fine moduli space $\mathcal{M}_{PI}(Y/X)$ for perverse ideal sheaves on Y . An S -valued point of $\mathcal{M}_{PI}(Y/X)$ is a family of perverse ideal sheaves on Y over S . By Proposition 3.9, applying the functor $\mathbf{R}f_*$ gives a family of ideal sheaves on X over S . This gives a morphism

$$\mathcal{M}_{PI}(Y/X) \longrightarrow \mathcal{M}_I(X),$$

where the scheme on the right is the moduli space of ideal sheaves on X .

On the other hand, an S -valued point of the Hilbert scheme $\text{Hilb}(X)$ determines a family of ideal sheaves on X over S , so there is a morphism

$$\text{Hilb}(X) \longrightarrow \mathcal{M}_I(X),$$

which induces a bijection on closed points. I claim that $\text{P-Hilb}(Y)$ is represented by the fibre product

$$\begin{array}{ccc} \text{P-Hilb}(Y/X) & \longrightarrow & \mathcal{M}_{PI}(Y/X) \\ \downarrow & \square & \downarrow \\ \text{Hilb}(X) & \longrightarrow & \mathcal{M}_I(X). \end{array}$$

Indeed, an S -valued point of the functor $\text{P-Hilb}(Y/X)$ is a triangle of objects

$$\mathcal{F} \longrightarrow \mathcal{O}_{S \times Y} \longrightarrow \mathcal{E} \longrightarrow \mathcal{F}[1]$$

of $\text{D}(S \times Y)$, each of which is a family of perverse sheaves on Y over S . Applying the functor $\mathbf{R}f_{S*}$ gives a short exact sequence of sheaves on $S \times X$, each of which is a family of sheaves on X by Proposition 3.9. This defines an S -valued point of $\text{Hilb}(X)$. The family \mathcal{F} defines an S -valued point of $\mathcal{M}_{PI}(Y/X)$ so we get an S -valued point of the fibre product.

Conversely, an S -valued point of the fibre product gives a family of perverse ideal sheaves \mathcal{F} on Y over S together with a short exact sequence

$$0 \longrightarrow f_{S*}(\mathcal{F}) \longrightarrow \mathcal{O}_{S \times X} \longrightarrow \mathcal{G} \longrightarrow 0$$

of S -flat sheaves on $S \times X$. By Lemma 7.2 this uniquely determines a triangle

$$\mathcal{F} \xrightarrow{\alpha} \mathcal{O}_{S \times Y} \longrightarrow \mathcal{E} \longrightarrow \mathcal{F}[1]$$

in $D(S \times Y)$ with $\mathcal{G} = \mathbf{R}f_{S*}(\mathcal{E})$. This is an S -valued point of the functor $\text{P-Hilb}(Y/X)$ providing \mathcal{E} defines a family of perverse sheaves on Y over S .

Applying the functor $\mathbf{L}j_s^*$ and noting that, as in Proposition 5.1, any nonzero morphism $\mathcal{F}_s \rightarrow \mathcal{O}_Y$ is an injection in $\text{Per}(Y/X)$, it is enough to show that $\mathbf{L}j_s^*(\alpha)$ is nonzero for all $s \in S$. But if $\mathbf{L}j_s^*(\alpha)$ vanishes then

$$H^{-1}(\mathbf{L}j_s^*(\mathcal{E})) = H^0(\mathcal{F}_s),$$

so $\mathbf{R}f_*(\mathbf{L}j_s^*(\mathcal{E}))$ cannot be a sheaf. But \mathcal{G} is flat over S , so this contradicts the argument of Proposition 3.9. \square

Proof of Theorem 3.8. Let (γ) denote the numerical equivalence class of the ideal sheaf \mathcal{I}_y of a point $y \in Y$. I claim that the scheme $\text{P-Hilb}(Y/X; \gamma)$ is a fine moduli space for perverse point sheaves on Y .

An S -valued point of $\text{P-Hilb}(Y/X; \gamma)$ certainly determines a family of perverse point sheaves on Y over S . For the converse suppose \mathcal{E} is a family of perverse point sheaves on Y over S . The object $\mathcal{G} = \mathbf{R}f_{S*}(\mathcal{E})$ is a family of points on X over S and hence, up to a twist by a line bundle from S , is the structure sheaf of the graph of a morphism $S \rightarrow X$. Twisting \mathcal{E} by the pullback of a line bundle from S we may assume that there is a surjection $\delta: \mathcal{O}_{S \times X} \rightarrow \mathcal{G}$ whose kernel is flat over S .

By adjunction, there is a morphism $\beta: \mathcal{O}_{S \times Y} \rightarrow \mathcal{E}$ such that $\mathbf{R}f_{S*}(\beta) = \delta$. Forming a triangle

$$\mathcal{F} \longrightarrow \mathcal{O}_{S \times Y} \xrightarrow{\beta} \mathcal{E} \longrightarrow \mathcal{F}[1]$$

gives an S -valued point of $\text{P-Hilb}(Y/X; \gamma)$ providing \mathcal{F} is a family of perverse sheaves. Applying the functor $\mathbf{L}j_s^*$ it will be enough to check that $\mathbf{L}j_s^*(\beta)$ is nonzero for all $s \in S$. But if $\mathbf{L}j_s^*(\beta) = 0$ then

$$H^1(\mathbf{L}j_s^*(\mathcal{F})) = H^0(\mathcal{E}_s),$$

so $\mathbf{R}f_*(\mathbf{L}j_s^*(\mathcal{F}))$ cannot be a sheaf, and this contradicts the argument of Proposition 3.9, since $\mathbf{R}f_{S*}(\mathcal{F})$ is flat over S . This completes the proof. \square

References

1. A. Beilinson, J. Bernstein, P. Deligne, Faisceaux Pervers, *Astérisque* **100**, Soc. Math de France (1983)
2. A.I. Bondal, Representations of associative algebras and coherent sheaves, *Math. USSR Izv.* **34** (1990), 23–42
3. A.I. Bondal, D.O. Orlov, Semiorthogonal decomposition for algebraic varieties, Preprint (1995) math.AG 9506012
4. T. Bridgeland, Equivalences of triangulated categories and Fourier-Mukai transforms, *Bull. London Math. Soc.* **31** (1999), no. 1, 25–34, also math.AG 9809114
5. T. Bridgeland, A. King, M. Reid, The McKay correspondence as an equivalence of derived categories, *J. Amer. Math. Soc.* **14** (2001), no. 3, 535–554, also math.AG 9908027

6. T. Bridgeland, A. Maciocia, Fourier–Mukai transforms for K3 and elliptic fibrations, to appear in *J. Alg. Geom.*, also math.AG 9908022
7. H. Clemens, J. Kollár, S. Mori, Higher dimensional complex geometry, *Astérisque* **166**, Soc. Math de France (1988)
8. S.I. Gelfand, Yu.I. Manin, *Methods of Homological Algebra*, Springer-Verlag (1996)
9. R. Hartshorne, *Residues and Duality*, *Lect. Notes Math.* **20**, Springer-Verlag, Heidelberg (1966)
10. J. Kollár, Flops, *Nagoya Math. J.* **113** (1989) 15–36
11. S. Mukai, On the moduli space of bundles on K3 surfaces I, in: *Vector Bundles on Algebraic Varieties*, M.F. Atiyah et al., Oxford University Press (1987), 341–413
12. C.T. Simpson, Moduli of representations of the fundamental group of a smooth projective variety I, *Publ. Math. I.H.E.S.* **79** (1994) 47–129
13. J.-L. Verdier, Des catégories dérivées des catégories abéliennes, *Astérisque* **239**, Soc. Math de France (1996)