D-Modules Over Smooth Affine Varieties

Contents

Introduction

A D-module is simply a module over a ring of differential operators, or more generally, a quasi coherent sheaf of modules over a sheaf of differential operators. Given this, it is no surprise that they first appeared in algebraic approaches to differential equations. The basic idea, dating back to B. Malgrange, is to associate a module M over the ring of differential operators D to a system of differential equations with solutions in some commutative ring of functions $\mathcal{O}$. The solution space of such a system can then be interpreted as $\operatorname{Hom}_D(M, \mathcal{O})$. This area of study truly took off with with M. Kashiwara's master's thesis, "Algebraic study of systems of partial differential equations" [Kas97], in which he applied techniques such as this to the study of equations with analytic coefficients.

Around the same time, I. N. Bernstein was developing the theory of modules over the Weyl algebra for entirely different reasons. He was interested in answering an open question posed by I. M. Gelfand, namely "what is the meaning of the complex power f^s of a polynomial?" This can be stated more concretely in terms of extending certain holomorphic functions on the upper half plane to meromorphic functions on all of $\mathbb{C}$. The first solutions to this question involved Hironaka's resolutions of singularities, but Bernstein discovered a more elementary answer involving the Weyl algebra and Bernstein-Sato polynomials, or b-functions [Ber72]. We discuss a version of his lemma on b-functions in these notes.

It is somewhat remarkable that D-modules as a field of study emerged independently from two starkly different starting points. Perhaps it was an early indication of the theory's ubiquity. Today, D-modules remain interesting due to a plethora of robust applications in fields as broad as mathematical physics, differential equations, representation theory and even number theory .

Where We Are Headed

These notes are intended to provide a friendly introduction to D-modules. However, a modern account of this theory requires some highly technical language, particularly from derived algebraic geometry and homological algebra, which we intend to avoid. This poses a severe limitation, as many standard definitions hinge on the use derived categories. It therefore becomes necessary to limit the scope of the discussion to avoid the most technical language pervading the theory.

As suggested by the title of this essay, our primary setting will be that of smooth affine varieties. The "smooth" hypothesis guarantees the existence of a so-called local coordinate system $\{x_i, \partial_i\}$ on $D_X|_U$ while the "affine" hypothesis means these local coordinate systems actually exist globally. Our goal is to provide as detailed an introduction to the basic theory of D-modules as possible while adhering to these two main assumptions. Ideally, this essay will develop enough intuition and background to prepare the reader for a more thorough treatment complete with homological and derived language. Luckily, much can be accomplished even in this restricted setting, for D-modules are characterized by their local structure.

Section 1 deals entirely with rings and sheaves of differential operators. It aims to provide insight into what is gained from the smooth and affine hypotheses and what is lost when they are omitted, and is hence longer than strictly necessary for the other sections. Section 2 introduces D-modules, focusing primarily on the case of the Weyl algebra. We discuss good filtrations, dimension, holonomy, and prove special cases of Bernstein's inequality and the lemma on b-functions. Section 3 is broadly about functoriality. It discusses pushforwards, pullbacks, and the celebrated Kashiwara equivalence.

We drop the affine hypothesis to discuss certain elementary results concerning D-modules over general smooth varieties where appropriate, but this is not our primary focus. Our approach is to first carefully develop the theory in the affine setting and then extend to the global case. Section 3 is the only exception, where we instead make global definitions before specializing to the affine case. This section deals with pushing forward and pulling back along a morphism $\varphi : X \longrightarrow Y$ of varieties, and there is substantial value in developing the vocabulary for such things in the global setting. We will, nonetheless, adopt two more major limitations to avoid derived functors; namely, we define direct images only for right D_X-modules along closed embeddings $\iota : X \hookrightarrow Y$.

Discussion of References

As this essay is entirely expository, all of the material presented here can be found in some form elsewhere. We focus exclusively on D-modules over smooth varieties with a heavy emphasis on the affine case. Such content lives somewhere between Coutinho's exposition on modules over the Weyl algebra [Cou95] and the book by Hotta, Takeuchi, and Tanisaki [HTT08], and the influence of both books is felt strongly throughout. The first chapter deals exclusively with rings and sheaves of differential operators, material largely inspired by online notes from Jeffries [Jef20] and from Ginzburg [Gin98]. The latter source proved invaluable for our discussion on good filtrations as well. Lecture notes by [Sch19] helped the author bridge the gap between the affine and non-affine cases, particularly in Sections 2.4, 3.2 and 3.3. The books [MR01] by McConnell and Robson and [BH93] by Bruns and Herzog were also helpful throughout Section 1.

1 Differential Operators

One must understand fields before one can define vectors spaces, and similarly one must understand the ring of differential operators before one can study D-modules. In this section, we do exactly that. We first define the ring of differential operators relative to an arbitrary ring homomorphism $A \longrightarrow R$ and discuss some of its basic properties before focusing on the case where A is a field and R a polynomial ring with coefficients in A. This latter object will provide a more explicit setting and will motivate arguments in the general case. We discuss several other examples of rings of differential operators, including in the case when R is not regular, and conclude this section by defining the sheaf of differential operators over a smooth variety.

It is worth noting that there are several equivalent ways to define the ring of differential operators in characteristic zero. We discuss three such definitions in the case of a polynomial ring over a field and show that they are equivalent when $\mathrm{char}(K) = 0$. These definitions will no longer coincide when $\mathrm{char}(K) > 0$, however.

1.1 The Ring of Differential Operators over an Arbitrary Ring

Let $A \longrightarrow R$ be a map of rings and let M and N be two R-modules. We may identify R with a subring of $\mathrm{End}_R(M)$ via the map which sends an element $f \in R$ to the R-linear map $\hat{f} : m \mapsto f \cdot m$ on M. We denote the image of $f \in R$ in $\mathrm{End}_R(M)$ by $\hat{f}_M$ when there is risk of confusing the domain of $\hat{f}$ with some other module. Given a morphism $\alpha \in \mathrm{Hom}_R(M, N)$, we abuse notation and write $[\alpha, \hat{f}]$ to mean $\alpha \circ \hat{f}_M - \hat{f}_N \circ \alpha$.

Definition 1.1. With A, R, M and N as above, we inductively define the collection of differential operators of order $k \in \mathbb{Z}$, denoted $D^k_{R/A}(M, N)$, as follows:

- $D^k_{R/A}(M, N) = 0$ when $k < 0$
- $D^k_{R/A}(M, N) = \left\{ \alpha \in \mathrm{Hom}_A(M, N) \;\middle|\; \left[\alpha, \hat{f}\right] \in D^{k-1}_{R/A}(M, N) \text{ for all } f \in R \right\}$ when $k \geq 0$.

We set $D_{R/A}(M, N) = \bigcup_{k \in \mathbb{Z}} D^k_{R/A}(M, N)$.

This is sometimes known as the "coordinate-free" approach to differential operators, and was first introduced by Grothendieck.

Remark 1.2. It is worth noting that $\alpha \in D_{R/A}(M, N)$ satisfies $[\alpha, \hat{f}] = 0 \in D^{-1}_{R/A}(M, N)$ exactly when α is R-linear, hence $D^0_{R/A}(M, N) = \mathrm{Hom}_R(M, N)$. Many sources, [Gin98] and [Ber] for instance, simply define $D^0_{R/A}(M, N) = \mathrm{Hom}_R(M, N)$ and proceed inductively from there.

Example 1.3. As a first example, suppose K is a field and R is a module finite K-algebra. Once we fix a basis for R, for any $f \in R$ the operator $\hat{f}$ is simply the diagonal matrix fI, where I is the identity matrix. Any other map $P \in \mathrm{Hom}_K(R, R)$, interpreted as a matrix, then satisfies

$$P \circ \hat{f} = P \cdot fI = fI \cdot P = \hat{f} \circ P,$$

hence $[P, \hat{f}] = 0$ and $P \in D^0_R$. It then follows that $D_R = \mathrm{Hom}_K(R, R)$.

We will see far more interesting examples later in section 1.2 and 1.3, but first we lay out some of the basis structure of rings of differential operators in general. The following lemma is elementary but nonetheless quite important:

Lemma 1.4. For each $k \in \mathbb{Z}$ we have an inclusion $D^{k-1}_{R/A}(M, N) \subseteq D^{k}_{R/A}(M, N)$. Furthermore, $D^{k}_{R/A}(M, N)$ is a left R-module under the action $f\alpha \mapsto \hat{f} \circ \alpha$ and a right R-module under the action $\alpha f \mapsto \alpha \circ \hat{f}$. This particularly implies that $R_{R/A}(M, N)$ is a left and right R-module under these same actions.

Proof: Both claims are proved inductively. The first is clear: the base case follows from the simple fact that $D^{-1}_{R/A}(M, N) = 0 \subseteq D^{0}_{R/A}(M, N)$, and if $\alpha \in D^{k-1}_{R/A}(M, N)$ then $[\alpha, \hat{f}] \in D^{k-2}_{R/A}(M, N)$ for any $f \in R$ by definition. The inductive hypothesis then implies that $[\alpha, \hat{f}] \in D^{k-1}_{R/A}(M, N)$, and hence $\alpha \in D^{k}_{R/A}(M, N)$.

For the second claim, note first that $\mathrm{Hom}_A(M, N)$ is an R-module by maps $R \longrightarrow \mathrm{Hom}_R(M, N) \longrightarrow \mathrm{Hom}_A(M, N)$, and since $D^{k}_{R/A}(M, N) \subseteq \mathrm{Hom}_A(M, N)$, it suffices to show that $D^{k}_{R/A}(M, N)$ is closed under addition and multiplication by R. Our base case is done by Remark 1.2: $D^{0}_{R/A}(M, N) = \mathrm{Hom}_R(M, R)$. Suppose then that $D^{m}_{R/A}(M, N)$ is a left R-module for each $m < k$ and note that for any two $f, g \in R$ the associated module endomorphisms commute by the commutativity of R, i.e. $\hat{f}\hat{g} = \hat{g}\hat{f}$. Fix $\alpha, \beta \in D^{k}_{R/A}(M, N)$ and $a, b \in R$. For any other $f \in R$ we have

$$\begin{aligned} [\hat{a}\alpha + \hat{b}\beta, \hat{f}] &= (\hat{a}\alpha + \hat{b}\beta)\hat{f} - \hat{f}(\hat{a}\alpha + \hat{b}\beta) \\ &= \hat{a}\alpha\hat{f} - \hat{a}\hat{f}\alpha + \hat{b}\beta\hat{f} - \hat{b}\hat{f}\beta \\ &= \hat{a}[\alpha, \hat{f}] + \hat{b}[\beta, \hat{f}]. \end{aligned}$$

Both $\hat{a}[\alpha, \hat{f}]$ and $\hat{b}[\beta, \hat{f}]$ are elements of the left R-module $D^{k-1}_{R/A}(M, N)$, hence so is their sum. The proof that $D^{k}_{R/A}(M, N)$ is a right R-module is similar. $\square$

Notation 1.5.

- We write $D_{R/A}(M)$ for $D_{R/A}(M, M)$ when $M = N$. As we shall see in Corollay 1.8, $D_{R/A}(M)$ is a ring under pointwise-addition and composition and is called the *ring of differential operators over M*. Given two operators $\alpha, \beta \in D_{R/A}(M)$ we often drop the composition symbol and write $\alpha\beta$ to mean $\alpha \circ \beta$.

- When $R = M = N$, we simply write $D_{R/A}$, or when there is no risk of ambiguity, D_R.

We will be primarily interested in D_R for a K-algebra R and will often write simply D_R. Some authors use this to denote the ring of differential operators with respect to the unique map $\mathbb{Z} \longrightarrow R$, but we never consider this case.

It will be useful to establish some basic commutator relations. These have nothing to do with differential operators but will used extensively in later sections, often without comment.

Proposition 1.6. Let A be a (not necessarily commutative) ring, M a left A-module and $\alpha, \beta, \gamma \in \mathrm{End}_A(M)$ A-linear maps on M. Then

(a) $[\alpha, \beta + \gamma] = [\alpha, \beta] + [\alpha, \gamma]$ and $[\alpha + \beta, \gamma] = [\alpha, \gamma] + [\beta, \gamma]$

(b) $[\hat{f}\alpha, \beta] = [\alpha, \hat{f}\beta] = \hat{f}[\alpha, \beta]$ for $f \in A$

(c) $[\alpha, \beta] = -[\beta, \alpha]$

(d) $[\alpha\beta, \gamma] = \alpha[\beta, \gamma] + [\alpha, \gamma]\beta$ and $[\alpha, \beta\gamma] = [\alpha, \beta]\gamma + \beta[\alpha, \gamma]$.

(e) $[\alpha, [\beta, \gamma]] + [\beta, [\gamma, \alpha]] + [\gamma, [\alpha, \beta]] = 0$ (Jacobi identity).

Proof: These are all straightforward computations. Identity (e) is perhaps slightly confusing, the left hand side is

$$\alpha(\beta\gamma - \gamma\beta) - (\beta\gamma - \gamma\beta)\alpha + \beta(\gamma\alpha - \alpha\gamma) + (\gamma\alpha - \alpha\gamma)\beta + \gamma(\alpha\beta - \beta\alpha) - (\alpha\beta - \beta\alpha)\gamma,$$

and all terms cancel once the expression is fully expanded. $\square$

1.1.1 Order of Differential Operators

Fix a commutative ring map $A \longrightarrow R$. A differential operator $D \in D_{R/A}(M)$ is said to be of *order* k if $D \in D^k_{R/A}(M)$ but $D \notin D^{k-1}_{R/A}(M)$. As the operator 0 is contained in $D^k_{R/A}$ for every $k \in \mathbb{Z}$, we say the order of 0 is $-\infty$. Here, we describe how order interacts with composition, addition, and commutation. Throughout this section $A \longrightarrow R$ is a map of commutative rings and M is an R-module.

Proposition 1.7. Suppose $\alpha \in D^m_{R/A}(M)$ and $\beta \in D^n_{R/A}(M)$. The following hold:

(a) $\alpha + \beta \in D^d_{R/A}(M)$ where $d = \max\{m, n\}$

(b) $\alpha\beta \in D^{m+n}_{R/A}(M)$

(c) $[\alpha, \beta] \in D^{m+n-1}_{R/A}(M)$.

Proof: Part (a) follows immediately from Lemma 1.4. We prove (b) and (c) simultaneously by induction on $m+n$. The base case is clear, for when $m+n=0$ we have $\alpha\beta \in \mathrm{Hom}_R(R,R)$. Suppose then that both (b) and (c) hold for $m+n<k$ for some positive integer k. Fix $f \in R$ and let $m+n=k$. By the inductive hypothesis we then have that $\alpha[\beta,\hat{f}]$ and $[\alpha,\hat{f}]\beta$ are in $D^{m+n-1}_{R/A}(M)$, and hence

$$[\alpha\beta, \hat{f}] = \alpha[\beta, \hat{f}] + [\alpha, \hat{f}]\beta \in D^{m+n-1}_{R/A}(M)$$

by Proposition 1.6 (d). This proves (b).

Rearranging the terms of the Jacobi identity, we have that

$$\big[[\alpha,\beta],\hat{f}\big] = \big[\alpha,[\beta,\hat{f}]\big] + \big[\beta,[\hat{f},\alpha]\big].$$

The inductive hypothesis tells us that the rightmost terms are elements of $D^{m+n-2}_{R/A}(M)$, hence so is $\big[[\alpha,\beta],\hat{f}\big]$. This proves (c). $\square$

This proposition yields some basic facts regarding the structure of $D_{R/A}(M)$.

Corollary 1.8. Let $A \longrightarrow R$ be a map of commutative rings. Then $D_{R/A}(M)$ is a ring and the graded ring

$$S_{R/A}(M) := \bigoplus_{k \in \mathbb{N}} S^k_{R/A}(M); \qquad S^k_{R/A}(M) = D^k_{R/A}(M)/D^{k-1}_{R/A}(M)$$

is commutative. We call $S_{R/A}(M)$ the *graded ring associated to* $D_{R/A}(M)$ and discuss it further in Section 2.

Proof: For any two $\alpha, \beta \in D_{R/A}(M)$, $\alpha\beta \in D_{R/A}(M)$ by Proposition 1.7 (b), hence $D_{R/A}(M)$ is a subring of $\mathrm{End}_A(M)$.

We identify $S^k_{R/A}(M)$ with its image under inclusion $S^k_{R/A}(M) \longrightarrow S_{R/A}(M)$ and let $\overline{\alpha}$ denote the image of $\alpha \in D^k_{R/A}(M)$ in $S^k_{R/A}(M)$. For $\alpha \in D^m_{R/A}(M)$ and $\beta \in D^n_{R/A}(M)$, we have $[\alpha, \beta] \in D^{m+n-1}_{R/A}(M)$ by Proposition 1.7 (c), hence $\overline{\alpha}\overline{\beta} - \overline{\beta}\overline{\alpha} = \overline{[\alpha, \beta]} = 0$. Since every element of $S_{R/A}(M)$ can be written as a sum of finitely many $\overline{\alpha}$, we are done. $\square$

1.1.2 Derivations

As of yet there has been no reason to restrict our generality, but now, we focus our attention exclusively on rings of differential operators of the form $D_{R/A}$. We already understand operators of order 0; since $D^0_{R/A} = \mathrm{Hom}_R(R, R) \cong R$, they're simply the operators of the form $\hat{f}$ for some $f \in R$. In this section we seek to understand the operators of order 1 as well, i.e. the R-module $D^1_{R/A}$.

Recall that an A-derivation of R is an A-linear map $d : R \longrightarrow R$ such that $d(ab) = ad(b) + d(a)b$ for all $a, b \in R$. Note that $d(1) = d(1 \cdot 1) = d(1) - d(1) = 0$. Further notice that for any derivation $d \in \mathrm{Der}_A(R)$ and $f, r \in R$,

$$[d, \hat{f}](r) = d\left(\hat{f}(r)\right) - \hat{f}\left(d(r)\right) = d(fr) - fd(r) = d(f)r.$$

This means that $[d, \hat{f}]$ is simply $\widehat{d(f)} \in D^0_{R/A}$ as a map on R, hence we have an inclusion $\iota : \mathrm{Der}_A(R) \hookrightarrow D^1_{R/A}$.

Let's now consider an arbitrary element $\alpha \in D^1_{R/A}$. The map $\alpha' = \alpha - \widehat{\alpha(1)}$ is also an order 1 operator by Lemma 1.4; in fact, it's a derivation. Indeed, it is A-linear by virtue of its membership to $D^1_{R/A}$ and for any $r, s \in R$ we have

$$\alpha'(rs) = \alpha'\hat{r}(s) = (\hat{r}\alpha')(s) + \widehat{\alpha'(r)}(s) = r\alpha'(s) + \alpha'(r)s$$

since $[\alpha', \hat{r}] = \alpha'(r)$.

Consider then the map $\varphi : D^1_{R/A} \longrightarrow \mathrm{Der}_A(R)$ defined $\varphi(\alpha) = \alpha - \widehat{\alpha(1)}$. It is A-linear, and since $\alpha(1) = 0$ for any derivation α, $\varphi \circ \iota$ is the identity on $\mathrm{Der}_A(R)$. This means the short exact sequence

$$0 \longrightarrow \ker\varphi \longrightarrow D^1_{R/A} \xrightarrow{\varphi} \mathrm{Der}_R(A) \longrightarrow 0$$

splits, giving us an isomorphism $D^1_{R/A} \cong \ker\varphi \oplus \mathrm{Der}_R(A)$. However, $\varphi(\alpha) = 0$ precisely when $\alpha = \widehat{\alpha(1)}$, i.e. when $\alpha \in D^0_{R/A} \cong \mathrm{Hom}_R(R, R) \cong R$. The results of this discussion are summarized in the proposition below.

Proposition 1.9. Let $A \longrightarrow R$ be a map of commutative rings. Then $D^1_{R/A} \cong R \oplus \mathrm{Der}_A(R)$ as A-modules via the map which sends $(f, d) \in R \oplus \mathrm{Der}_A(R)$ to $\hat{f} + d$. $\square$

It is important to note that there is a more functorial way to define derivations. Given an A-algebra R, we first define the multiplication map $R \otimes_A R \longrightarrow R$ given by $x \otimes y \mapsto xy$. The kernel of this map is denoted $\Delta_{R/A}$ and is generated by elements of the form $r \otimes 1 - 1 \otimes r$:

$$\Delta_{R/A} = \langle \{r \otimes 1 - 1 \otimes r \mid r \in R\} \rangle = \ker(R \otimes_A R \xrightarrow{mult} R). \tag{1}$$

We use this to define the module of Kähler differentials.

Definition 1.10. Let R be an A-algebra. The module of A-linear Kähler differentials is

$$\Omega_{R/A} = \Delta_{R/A}/\Delta^2_{R/A}.$$

It comes equipped with a derivation $d : R \longrightarrow \Omega_{R/A}$ called the *universal derivation*:

$$d(r) = r \otimes 1 - 1 \otimes r \; + \; \Delta^2_{R/A}.$$

Hartshorne defines $\Omega_{R/A}$ to be the R-module, unique up to isomorphism, equipped with an A-derivation $d : R \longrightarrow \Omega_{R/A}$ such that for any other A-derivation $d' : R \longrightarrow M$ there exists a unique R-module map $f : \Omega_{R/A} \longrightarrow M$ with $d' = f \circ d$. This is equivalent to the definition given above. Defining $\Omega_{R/A}$ via a universal property does make the following characterization of $\mathrm{Der}_A(R)$ immediately evident, however:

Proposition 1.11. Let M be an R-module. There exists an isomorphism of R-modules

$$\mathrm{Hom}_R(\Omega_{R/A}, M) \cong \mathrm{Der}_A(M)$$

given by composing a map $f : \Omega_{R/A} \longrightarrow M$ with the universal derivation $d : R \longrightarrow \Omega_{R/A}$. Hence the functor $M \mapsto \mathrm{Der}_A(M)$ is represented by $\Omega_{R/A}$.

In other words, $\Omega_{R/A}$ represents the functor $\mathrm{Der}_A(-)$.

1.1.3 Derivation Examples

Proposition 1.9 tells us that to understand $D^1_{R/A}$ it suffices to understand $\mathrm{Der}_A(R)$. Here, we explicitly describe the module $\mathrm{Der}_A(R)$ for specific rings R.

Example 1.12. Let K be a field of characteristic zero and $R = K[x_1, ..., x_n]$ be a polynomial ring over K. By the product rule, the K-linear maps ∂_{x_i} $(1 \leq i \leq n)$ each of which sends a polynomial f to its partial derivative in x_i are derivations. Any other derivation $\alpha \in \mathrm{Der}_K(R)$ satisfies

$$\alpha(x_i^k) = kx_i^{k-1}\alpha(x_i) = \partial_{x_i}(x_i^{k_i})\alpha(x_i).$$

This means that for a monomial $x_1^{k_1}...x_n^{k_n}$ we have

$$\begin{aligned}
\alpha\left(x_1^{k_1}...x_n^{k_n}\right) &= \alpha(x_1^{k_1})x_2^{k_2}...x_n^{k_n} \ + \ x_1^{k_1}\alpha(x_2^{k_2}...x_n^{k_n}) \\
&= \alpha(x_1)\partial_{x_1}(x_1^{k_1}...x_n^{k_n}) \ + \ x_1^{k_1}\left(\alpha(x_2^{k_2})x_3^{k_3}...x_n^{k_n} \ + \ x_2^{k_2}\alpha(x_3^{k_3}...x_n^{k_n})\right) \\
&\ \vdots \\
&= \alpha(x_1)\partial_{x_1}(x_1^{k_1}...x_n^{k_n}) \ + \ ... \ + \ \alpha(x_n)\partial_{x_1^{k_1}...x_n^{k_n}}.
\end{aligned}$$

Since monomials form a basis over K for R, we get that $\alpha = \alpha(x_1)\partial_{x_1} + ... + \alpha(x_n)\partial_{x_n}$. Hence $\{\partial_{x_1}, ..., \partial_{x_n}\}$ generates $\mathrm{Der}_K(R)$ as a R-module. In particular, $\mathrm{Der}_K(R)$ is a free-module over R of rank n.

Example 1.13. As before, let K be a field of characteristic zero. Consider the ring $R = K[t^2, t^3]$, noting that $R \cong K[x, y]/J$ for $J = (y^2 - x^3)$ via the map $x \mapsto t^2$ and $y \mapsto t^3$. As we will see, $\mathrm{Der}_K(R)$ is generated by $t\partial_t$ and $t^2\partial_t$.

First consider the derivations $D_1 = 2y\partial_x + 3x^2\partial_y$ and $D_2 = 3y\partial_y + 2x\partial_x$ on $K[x, y]$. These are also derivations on $K[x, y]/J$ since $D_1(J), D_2(J) \subseteq J$, and in fact, we will show they generate $\mathrm{Der}_K(K[x, y]/J)$ as a $K[x, y]/J$-module. Any other derivation α on $K[x, y]/J$ can be written as $\alpha = f_1\partial_x + f_2\partial_y$ by the previous example with the extra requirement that $\alpha(J) \subseteq J$. This is equivalent to the condition

$$-3x^2 f_1 \ + \ 2yf_2 = u(y^2 - x^3) \tag{2}$$

for some polynomial $u \in K[x, y]$. Notice that f_1 cannot have a constant term, if it did, the LHS of equation (2) would have a nonzero x^2 summand while the RHS would not. This means f_1 may have only terms of degree 1 or higher, hence we may write $f_1 = 2(yg + xh)$ for some $g, h \in K[x, y]$. Plugging this into equation (2) and rearranging yields

$$2yf_2 = u(y^2 - x^3) + 6x^2yg + 6x^3h$$

and substituting $u' = u - 6(y^2 - x^3)h$ gives

$$2yf_2 = u'(y^2 - x^3) + 6x^2yg + 6x^3h + 6(y^2 - x^3)h = u'(y^2 - x^3) + 6x^2yg + 6y^2h.$$

The LHS of this equation is divisible by y so the RHS must be as well, implying $v = \frac{u'}{2y} \in K[x, y]$. Hence $f_2 = v(y^2 - x^3) + 3x^2g + 3yh$. We then get

$$\begin{aligned}\alpha = f_1\partial_x + f_2\partial_y &= 2(yg + xh)\partial_x + \left(v(y^2 - x^3) + 3x^2g + 3yh\right)\partial_y \\ &= g(2y\partial_x + 3x^2\partial_y) + h(2x\partial_x + 3y\partial_y) + v(y^2 - x^3)\partial_y \\ &= gD_1 + hD_2 + v(y^2 - x^3)\partial_y.\end{aligned}$$

Since $v(y^2 - x^3)\partial_y$ is the trivial derivation on $K[x, y]/J$, the above shows that α is in the $K[x, y]/J$-span of D_1 and D_2. Finally, for an arbitrary $f \in R$ we have

$$t\partial_t(f) = t \cdot \frac{\partial f}{\partial x}\frac{dx}{dt} + t \cdot \frac{\partial f}{\partial y}\frac{dy}{dt} = 2t^2\frac{\partial f}{\partial x} + 3t^3\frac{\partial f}{\partial_y} = (2x\partial_x + 3y\partial_y)(f) = D_2(f)$$

and

$$t^2\partial_t(f) = t^2 \cdot \frac{\partial f}{\partial x}\frac{dx}{dt} + t^2 \cdot \frac{\partial f}{\partial y}\frac{dy}{dt} = 2t^3\frac{\partial f}{\partial x} + 3t^4\frac{\partial f}{\partial y} = (2y\partial_x + 3x^2\partial_y)(f) = D_1(f)$$

by the chain rule.

1.2 The Weyl Algebras

Throughout this section $A = K$ and $R = K[x_1, ..., x_n]$, where K is a field. We call the ring D_R the n^{th} *Weyl Algebra*, and it was one of the earliest rings of differential operators to be studied in detail. It first appeared as Dirac's *quantum algebra*, which consists of polynomial expressions in variables p and q subject to the relation $pq - qp = 1$. Weyl algebras admit tractable, explicit descriptions in terms of generators and relations and thereby serve as a fantastic source of examples. They also provide a good starting point for newcomers seeking to develop intuition.

Our first aim in this section is to show the three main presentations of the n$^{\text{th}}$ Weyl algebra are equivalent.

Theorem 1.14. *(Definition) Let K be a field of characteristic 0 and let $R = K[x_1, ..., x_n]$. The following are isomorphic modules.*

- *The K-subalgebra $A_n(K) \subseteq \mathrm{End}_K(R)$ generated by the maps $\hat{x}_i$ and $\partial_{x_i} = \frac{\partial}{\partial x_i}$. We will often omit the K from this notation when there is no risk of ambiguity and simply write A_n.*
- *The K-algebra D_n defined to be the free K-algebra in the $2n$-variables $y_1, ..., y_{2n}$ modulo the ideal J, where multiplication is given by concatenation on monomials and J is generated by all the elements of the form $[y_{i+n}, y_i] - 1$ for $1 \leq i \leq n$ or $[y_a, y_b]$ for $a \not\equiv b \mod n$, $1 \leq a, b \leq 2n$.*

- *The ring of differential operators D_R.*

Before we prove this, we need to understand some basic facts about the module A_n.

Lemma 1.15. The generators of A_n satisfy the following relations:

$$[\partial_{x_i}, \hat{x}_j] = \delta_{ij}, \qquad [\partial_{x_i}, \partial_{x_j}] = [\hat{x}_i, \hat{x}_j] = 0$$

where δ_{ij} is the Kronecker delta function. Furthermore, for $f \in R$,

$$[\partial_{x_i}, \hat{f}] = \widehat{\frac{\partial f}{\partial x_i}}.$$

Proof: For any polynomial f (and more generally, any differentiable function) we have

$$\partial_{x_i}\hat{x}_j(f) = \partial_{x_i}(x_j \cdot f) = \partial_{x_i}(x_j) \cdot f \ + \ x_j \cdot \partial_{x_i}(f)$$

from the product. Since $\partial_{x_i}(x_j) = \delta_{ij}$ and $x_j \cdot \partial_{x_i}(f) = \hat{x}_j\partial_{x_i}(f)$, rearranging the above yields the first relation.

Differentiation is K-linear, so it suffices to prove $\partial_{x_i}\partial_{x_j}(f) = \partial_{x_j}\partial_{x_j}(f)$ for a monomial f. This is clear from the power rule. The fact $[\hat{x}_i, \hat{x}_j] = 0$ is a consequence of the commutativity of x_i and x_j in R.

Finally, for the final property it once again suffices to prove $[\partial_{x_i}, \hat{f}] = \frac{\partial f}{\partial_{x_i}}$ for monic monomials. We first show it holds for $f = x_i^m$. The relation $[\partial_{x_i}, \hat{x}_1] = 1$ serves as the base case, so suppose it holds for all $m < k$. Then

$$\partial_{x_i}\hat{x}_i^k = (\partial_{x_i}\hat{x}_i)\hat{x}_i^{k-1} = (1 + \hat{x}_i\partial_{x_i})\hat{x}_i^{k-1} = \hat{x}_i^{k-1} + \hat{x}_i\partial_{x_i}\hat{x}_i^{k-1}.$$

The inductive hypothesis implies $\partial_{x_i}\hat{x}_i^{k-1} = (k-1)\hat{x}_i^{k-1} + \hat{x}_i^{k-1}\partial_{x_i}$, so after rearranging the above and combining like terms we have exactly that $[\partial_{x_i}, \hat{x}_i^k] = k\hat{x}_i^{k-1}$.

For an arbitrary monic monomial $x_1^{m_1}...x_n^{m_n}$ we have that

$$[\partial_{x_i}, \hat{x}_1^{m_1}...\hat{x}_n^{m_n}] = \hat{x}_1^{m_1}...\hat{x}_{i-1}^{m_{i-1}}[\partial_{x_i}, x_i^{m_i}]\hat{x}_{i+1}^{m+1}...\hat{x}_n^{m_n}$$

by repeated use of Proposition 1.6 (d). This reduces to

$$[\partial_{x_i}, \hat{x}_1^{m_1}...\hat{x}_n^{m_n}] = m_i \cdot \hat{x}_1^{m_1}...\hat{x}_i^{m_i-1}...\hat{x}_n^{m_n}$$

by what we have already proven. □

Remark 1.16. It is worth saying a few words about our choice of notation. Most authors simply write "f" to refer interchangeably to $f \in R$ and its image in $D_{R/A}(M)$. This is reasonable, especially since the R-action on $D_{R/A}(M)$ is given by the inclusion $R \hookrightarrow D_{R/A}(M)$. Nonetheless, we prefer to distinguish between an element $f \in R$ and its image in $D_{R/A}(M)$ due to the notational similarity between $\partial_{x_i}\hat{f}$ and $\partial_{x_i}(f)$. There exist abundant opportunities for confusion; for example, $\partial_x(x) = 1 \in K[x]$ whereas $\partial_x\hat{x} = 1 + \hat{x}\partial_x \neq 1 \in A_1$.

We now construct a basis for the Weyl algebra, a basis known as the *canonical basis*.

Lemma 1.17. The set $\mathbf{B} = \{\hat{x}^\alpha\partial^\beta \mid \alpha, \beta \in \mathbb{N}^n\}$ is a basis for A_n as a K-vector space. By $\hat{x}^\alpha$ we mean the operator $\hat{x}_1^{\alpha_1} \cdot ... \cdot \hat{x}_n^{\alpha_n}$, and the degree of this monomial is the length of α defined $|\alpha| = \alpha_1 + ... + \alpha_n$.

Proof: By definition, A_n is generated by monomials in ∂_{x_i} and $\hat{x}_j$ for i and j ranging between 1 and n. Using the fact that $\partial_{x_i}\hat{x}_i - \hat{x}_i\partial_{x_i} = \widehat{\frac{\partial f}{\partial x_i}}$ from Lemma 1.15 we can move all $\hat{x}_j$ terms to the left of all ∂_i terms, so it is clear that $\mathbf{B}$ spans A_n.

We now show that $\mathbf{B}$ is linearly independent. Suppose that

$$D = \sum_{i=1}^{m} c_i \hat{x}^{\alpha_i} \partial^{\beta_i}.$$

We call this summation the *canonical form* of $D \in A_n$ and show that $D = 0$ if and only if $c_i = 0$ for each $1 \leq i \leq m$. Assume without loss of generality that $c_i \neq 0$ for all $1 \leq i \leq m$ and $(\alpha_i, \beta_j) = (\alpha_j, \beta_j)$ if and only if $i = j$; that is, make m as small as possible. Let β_ℓ be the multi-index such that $|\beta_\ell| = \min\{|\beta_1|, ..., |\beta_m|\}$. By repeated use of the power law we get that

$$\partial^{\beta_\ell}(x^{\beta_\ell}) = \beta_\ell! \neq 0$$

where $\beta! = \beta_1! \cdot ... \cdot \beta_n!$ for $\beta \in \mathbb{N}^n$, but that $\partial^{\beta_i}(x^{\beta_\ell}) = 0$ for all $|\beta_i| > |\beta_\ell|$. It is possible that ∂^{β_ℓ} appears multiple times in the above summation. For simplicity, set $\lambda = \beta_\ell!$ and let $\{\alpha_1', ..., \alpha_k'\}$ be the (necessarily distinct) multi-indices such that $\hat{x}^{\alpha_i'}\partial^{\beta_\ell}$ appears with nonzero coefficient in the canonical form of D. Likewise let c_i' be the coefficient of $\hat{x}^{\alpha_i'}\partial^{\beta_\ell}$ appearing in the canonical form of D. Then

$$D(x^{\beta_\ell}) = \sum_{i=1}^{k} c_i' \hat{x}^{\alpha_i'} \partial^{\beta_\ell}(x^{\beta_\ell}) = \lambda \left(c_1' x^{\alpha_1'} + ... + c_k' x^{\alpha_k'} \right).$$

Since the α_i' are pairwise distinct, the above polynomial is nonzero and $D \neq 0$. We conclude that $D = 0$ if and only if $c_i = 0$ and we conclude that $\mathbf{B}$ is linearly independent over K. $\square$

To illuminate the details of the above proof, let's examine some examples of differential operators over a polynomial ring in canonical form.

Example 1.18. Consider the first Weyl algebra $D_{K[x]}$, which is generated by $\hat{x}$ and ∂. The following identities hold:

(a) $\partial^m \hat{x} = \hat{x}\partial^m + m \cdot \partial^{m-1}$ and

(b) $\partial^a \hat{x}^b = \sum_{j=0}^{d} j! \binom{a}{j} \binom{b}{j} \hat{x}^{b-j} \partial^{a-j}$.

These of course easily generalize to D_R by replacing $\hat{x}$ with $\hat{x}_i$ and ∂ with ∂_i. They are both proven via induction and liberal use of the fact that $[\partial, \hat{x}^b] = b\hat{x}^{b-1}$, but neither proof is particularly enlightening. It is perhaps more useful to see an explicit computation for low values of a and b:

$$\begin{aligned} \partial^2 \hat{x}^3 &= \partial \left(\partial \hat{x}^3 \right) \\ &= \partial \left(\hat{x}^3 \partial \ + \ 3\hat{x}^2 \right) \\ &= \hat{x}^3 \partial^2 \ + \ 6\hat{x}^2 \partial \ + \ 6\hat{x} \end{aligned}$$

and how (b) can be used to compute the canonical form of operators in larger Weyl algebras, for instance in

$D_{K[x,y]}$:

$$\begin{aligned}\partial_x\partial_y^2\hat{x}^3\hat{y}^2 &= \partial_x^2\hat{x}^3 \cdot \partial_6^2\hat{y}^2 \\ &= \left(\hat{x}^3\partial_x \;+\; 3\hat{x}^2\right)\left(\hat{y}^2\partial^2 \;+\; 4\hat{y}\partial_y \;+\; 2\right) \\ &= \hat{x}^3\hat{y}\partial_x\partial_y^2 \;+\; 3\hat{x}^2\hat{y}^2\partial_y^2 \;+\; 4\hat{x}^3\hat{y}\partial_x\partial_y \;+\; 12\hat{x}^2\hat{y}\partial_y \;+\; 2\hat{x}^3\partial_x \;+\; 6\hat{x}^2.\end{aligned}$$

In the general setting of $D_{R/A}$ where $A \longrightarrow R$ is an arbitrary map of rings, we have a notion of order. For the ring of differential operators over a polynomial ring, the existence of the canonical basis gives us something something better: a notion of degree. This doesn't give us a graded structure, but it does recover some of the properties of degree in a polynomial ring.

Let $D \in A_n$ be an operator in canonical form. The degree of D, denoted $\deg(D)$, is the length $|(\alpha, \beta)|$ of the largest multindex $(\alpha, \beta) \in \mathbb{N}^n \times \mathbb{N}^n$ such that $x^\alpha\partial^\beta$ appears with nonzero coefficient in the canonical form of D. The following proposition should be compared to Proposition 1.7, and due to its similarity the proof is omitted (Hint: as with many things in life, it suffices to check monomials).

Proposition 1.19 ([Cou95, Theorem 2.1.1.])**.** Let $D, D' \in A_n$ and assume $\mathrm{char}(K) = 0$.

(a) $\deg(DD') = \deg(D) + \deg(D')$

(b) $\deg(D + D') \leq \max\{\deg(D), \deg(D')\}$

(c) $\deg[D, D'] \leq \deg(D) + \deg(D') - 2$.

As $\deg(0) = -\infty$, an immediate corollary to part (a) of the above proposition is that A_n is a domain. We can also use the proposition to prove the following theorem:

Theorem 1.20. *The algebra A_n is simple.*

Proof: Let I be a nonzero two-sided ideal of A_n and suppose $D \in I$ is a nonzero operator. If $deg(D) = 0$, then $D \in K$ and $I = A_n$. If $\deg(D) = d > 0$, then there must be some summand $x^\alpha\partial^\beta$ with nonzero coefficient and for which either $\alpha \neq 0$ or $\alpha \neq 0$. In the former case, suppose the α_i component of α is nonzero. Then $[\partial_i, D] \neq 0$ and $\deg([\partial_i, D]) \leq d - 1$. Furthermore, since I is two-sided, $[\partial_i, D] \in I$. By replacing D with $[\partial_i, D]$ and repeating the above process, we can construct an element of degree 0 in I and hence conclude$I = A_n$. A similar argument in which we instead consider $[x_i, D]$ works in the case that $\beta \neq 0$. $\square$

Note that while A_n does not have any proper nontrivial two-sided ideals, it has many left and right ideals and is by no means a division ring. Furthermore, the kernel of any map of nontrivial unital rings must necessarily be a two-sided ideal, hence we have the following corollary.

Corollary 1.21. If $\phi : A_n \longrightarrow B$ is a map of unital rings then it is injective. $\square$

We are now ready to prove Theorem 1.14.

Proof: (Theorem 1.14) We first show $A_n \cong D_n$. Let $K\{y_1, ..., y_{2n}\}$ denote the free algebra over K in $2n$ variables with multiplicative given by concatenation of monomials and let $J \subseteq K\{y_1, ..., y_{2n}\}$ be the ideal generated by all the elements of the form $[y_{i+n}, y_i] - 1$ for $1 \leq i \leq n$ or $[y_a, y_b]$ for $a \not\equiv b \mod n$, $1 \leq a, b \leq 2n$. Note $D_n = K\{y_1, ..., y_{2n}\}/J$ by definition.

Define a map $\psi : A_n \longrightarrow D_n$ by setting $\psi(x^\alpha \partial^\beta) = y^{(\alpha,\beta)} + J$, noting that it suffices to define ψ on monomials in canonical form. A quick check shows that each of the relations on the generators of A_n given in Lemma 1.15 are preserved by ψ, so it is indeed a map of rings. Using the relations given by J, the same proof used in Lemma 1.17 can be used to show $\{y^{\alpha,\beta} + J\}_{\alpha,\beta \subseteq \mathbb{N}^n}$ is a basis for D_n, so it is clear that ψ is surjective. Furthermore, ψ is a map of unital rings and is therefore injective by Corollary 1.21. Hence ψ is an isomorphism.

We now wish to prove $A_n \cong D_R$. Denote by C_k the subset of A_n consisting of operators of degree at most k. We use the following two facts without proof:

(i) If $P \in D_R$ and $[P, \hat{x}_i] = 0$ for each $1 \le i \le n$, then $P \in R$ ([Cou95, Lemma 3.2.1]).

(ii) Let $P_1, ..., P_n \in C_{r-1}$ and assume that $[P_i, x_j] = [P_j, x_i]$ for all $1 \le i, j \le n$. Then there exists $Q \in C_r$ such that $P_i = [Q, x_i]$, for $i = 1, ..., n$ ([Cou95, Lemma 3.2.2]).

From Proposition 1.19 it is clear that $C_k \subseteq D_R^k$, so it suffices to prove the reverse inclusion. We proceed by induction. Proposition 1.9 gives us the base case $k = 1$. Suppose then that $D_R^r = C_r$ for all $0 \le r \le k-1$ and that $P \in D_R^k$. Let $P_i = [P, \hat{x}_i]$ and note that $P_i \in D_R^{k-1}$ by definition. Since $\hat{x}_i$ and $\hat{x}_j$ commute for all $1 \le i, j \le n$ we have

$$[P_i, x_j] = [[P, x_i], x_j] = [[P, x_j], x_i] = [P_j, x_j]$$

by the Jacobi identity. By fact (ii) above, there exists some $Q \in C_k$ such that $[Q, x_i] = P_i$ for each $1 \le i \le n$ and hence $[Q - P, x_i] = 0$. Then $Q - P \in R$ by fact (i) above, so $P = Q + \hat{f}$ for some $f \in R$. This means $P \in C_k$, and we are done. $\square$

1.2.1 Difficulties in Prime Characteristic

Even at this early stage, when K is of positive characteristic $p > 0$ we can see major departures from what we have shown. Consider $A_1 = K[x, \partial] \subseteq \mathrm{End}_K(K[x])$ for $K = \mathbb{F}_p$. Let k be any positive integer and consider the action ∂^p on $x^k \in K[x]$. If $k < p$, then $\partial^p(x^k) = 0$. If $k \ge p$, then at least one of the integers $k-p+1, k-p+2, ..., k-1, k$ is divisible by p, and hence

$$\partial^p(x^k) = k(k-1)(k-2)...(k-p+1)x^{k-p} = 0.$$

Since ∂^p is zero on a basis for $K[x]$, it is identically zero on all of $K[x]$. This means ∂ is a nilpotent element and hence A_1 is not a domain.

Now consider D_1, the free algebra in x and ∂ over K modulo the relation $[\partial, x] = 1$. In contrast to A_1, this ring is a domain since Proposition 1.19 still holds, so we no longer have $A_1 \cong D_1$. It is not clear that D_1 is preferable to A_1 however, for there is another major departure from the characteristic zero world: D_n is not simple. For example,

$$[\partial, x^p] = px^{p-1} = 0,$$

from which it follows that D_1 has a nontrivial center, a two-sided ideal.

Furthermore, in characteristic zero, not all operators can be written as R-linear combinations of compositions of derivations. Take for instance the operator $\alpha \in D_{R/\mathbb{F}_p}$ when $R = \mathbb{F}_p[x]$ defined

$$x^n \mapsto \begin{cases} \binom{n}{p} x^{n-p} & \text{if } n \geq p \\ 0 & \text{otherwise} \end{cases}.$$

In characteristic zero, this operator is simply $\frac{1}{p!}\partial^p$, but in characteristic $p > 0$ it cannot be written as the composition of smaller order operators.

To summarize, when working with rings of differential operators D_R, it is necessary to fix either the characteristic of K to be 0 or the choice of definition for D_R. In this document we do the former and consider only fields of characteristic 0.

1.3 Differential Operators on a Smooth Variety

It seems natural to ask whether there exist nice descriptions of D_R comparable to those given by Theorem 1.14 and Lemma 1.17 when R is "nearly a polynomial ring". When "nearly a polynomial ring" is interpreted to mean "a regular K-algebra of finite type", the answer turns out to be "yes". The regular hypothesis is quite necessary, as we shall see. Regular finitely-generated K-algebras are also precisely the local version of smooth algebraic varieties, which we introduce in the context of differential operators here.

1.3.1 Regular *K*-Algebras of Finite Type

Theorem 1.22. *Let R be a regular K-algebra of finite type. Then D^m_R is generated as an R-module by all products of up to m many K-derivations of R. In particular, D_R is generated by R and $\mathrm{Der}_K(R)$ as an R-module.*

Proof: The case in which R is a domain is handled by [MR01, Theorem 15.5.5]. Here is a rough outline of the ideas used. Suppose $L = \mathrm{Frac}(R)$ and $\{x_1, ..., x_n\}$ is a transcendence basis for L over K. One can pass to the polynomial ring $K[x_1, ..., x_n]$ and use the fact that $\mathrm{Der}_K(L) = \sum L \cdot \partial/\partial x_i$ to show that D_L is spanned by L and $\mathrm{Der}_K(L)$ by mimicking the proof of the case in which R is a polynomial ring. It then only remains to prove $D_R = \{\alpha \in D_L \mid \alpha(R) \subseteq R\}$.

The general case is given by [Muh88, Theorem 1.15]. Every regular ring is reduced, hence the intersection of all minimal primes in R is 0. The ring R can therefore be written as a product of domains by the Chinese Remainder Theorem. Muhasky uses the fact that $D_{(R_1 \times R_2)} \cong D_{R_1} \times D_{R_2}$ to conclude. □

Not only do derivations generate D_R, but each operator can be expressed in a way reminiscent of the canonical form for operators in the Weyl algebra (see Lemma 1.17). Namely, if R is a regular K-algebra of Krull dimension n, then any $P \in D^k_R$ can be written as a finite sum

$$P = \sum_{\alpha} \hat{f}_\alpha \partial^\alpha$$

where each $\alpha \in \mathbb{N}^n$, $f_\alpha \in R$ and $\{\partial_1, ..., \partial_n\}$ generate $\mathrm{Der}_K(R)$. This fact is slightly stronger than Theorem 1.22 however. See Theorem 1.26 below.

Let us examine two examples, one in which the hypotheses of Theorem 1.22 hold and one in which they do not.

Example 1.23. Let K be a field of characteristic zero and set $R = K[x, y]/(f)$ where $f = x^3 - x - y^2$. As the matrix

$$\begin{bmatrix} \frac{\partial f}{\partial x}(x_0, y_0) & \frac{\partial f}{\partial y}(x_0, y_0) \end{bmatrix} = \begin{bmatrix} 3x_0^2 - 1 & -2y_0 \end{bmatrix}$$

is rank 1 for all points $(x_0, y_0) \in K^2$ in the graph of f, R is easily seen to be regular by the Jacobian criterion. Hence, to understand D_R it suffices to understand the derivations on R.

It isn't terribly difficult to see that the set of derivations on R is given by

$$\operatorname{Der}_K(R) = \frac{\{\theta \in \operatorname{Der}_K(K[x, y]) \mid \theta((f)) \subseteq (f)\}}{(f) \operatorname{Der}_K(K[x, y])}.$$

We know $\operatorname{Der}_K(R)$ is a one-dimensional K-vector space since $\operatorname{Der}_K(R)$ is two-dimensional by example 1.12. It therefore suffices to find one derivation $\theta : K[x, y] \longrightarrow K[x, y]$ which fixes (f) to compute $\operatorname{Der}_K(R)$. Furthermore, since $\theta(f \cdot g) = f\theta(g) + g\theta(f)$, θ fixes (f) if and only if $\theta(f) \in (f)$, reducing our task of calculating $\operatorname{Der}_K(R)$ to finding a single derivation θ on $K[x, y]$ which sends f to a multiple of itself. But this is exceptionally easy; the derivation $\theta = \partial_x(f)\partial_y - \partial_y(f)\partial_x$ maps f to zero.

We conclude that $D_R = \bigoplus_{k=0}^{\infty} R \cdot \theta^k$ where $\theta = (3x^2 - 1)\partial_y + 2y\partial_x$.

Example 1.24. We return to the curve $f = y^2 - x^3$, which has a singularity at the origin. Let $R = K[t^2, t^3]$ and recall from Example 1.13 that $K[x, y]/(f) \cong K[t^2, t^3]$.

Consider the operator $\alpha = t\partial_t^2 - \partial_t$ in $D_{K[t]}$. Since $\alpha(t^2) = 0$ and $\alpha(t^3) = 3t^2$, $\alpha(R) \subseteq R$ and therefore $\alpha|_R \in D_R$. However, $\operatorname{Der}_K(R)$ is generated as a vector space by $t\partial_t$ and $t^2\partial$, and by considering these to be operators on $K[t]$ it is clear that α is outside the subring of $D_{K[t]}$ generated by $t^2\partial_t$ and $t\partial_t$. Therefore D_R is strictly larger than the ring generated by $\operatorname{Der}_K(R)$ and R, highlighting the need for the regular hypothesis in Theorem 1.22.

1.3.2 Smooth Varieties

We now define the sheaf of differential operators on a smooth variety, the primary setting of [HTT08]. The definitions given here are precisely those found in section 1.1 of [HTT08] contextualized within the discussion up to this point.

Definition 1.25. Let X be a smooth variety over a field K of characteristic zero and $\mathcal{O}_X$ be its structure sheaf. We denote by $\mathcal{E}nd_K\, \mathcal{O}_X$ the sheaf of K-linear endomorphisms of $\mathcal{O}_X$. We say that a section $\theta \in (\mathcal{E}nd_K\, \mathcal{O}_X)(X)$ is a *vector field on* X if $\theta(U) = \theta|_U$ is a K-derivation on $\mathcal{O}_X(U)$ for each open subset $U \subseteq X$. For any open subset $U \subseteq X$, the set of vector fields on U is denoted $\Theta(U)$. Then $\Theta(U)$ is an $\mathcal{O}_X(U)$-module, and the assignment $U \mapsto \Theta(U)$ is a sheaf of $\mathcal{O}_X$-modules. We denote this sheaf by Θ_X and note that when X is affine, $\Theta_X \cong \widetilde{\operatorname{Der}_K}(\mathcal{O}_X(X))$.

We then have the following theorem.

Theorem 1.26. *Let X be a smooth algebraic variety of dimension n over an algebraically closed field K. Then for each point $p \in X$, there exists an affine open neighborhood V of p, regular functions $x_i \in K[V] = \mathcal{O}_X(V)$, and vector fields $\partial_i \in \Theta_X(V)$ for $1 \leq i \leq n$ satisfying the conditions*

$$\begin{cases} [\partial_i, \partial_j] = 0, \quad \partial_i(x_j) = \delta_{ij} \ (1 \leq i, j \leq n) \\ \Theta_V = \bigoplus_{i=1}^{n} \mathcal{O}_V \partial_i \end{cases}.$$

Moreover, we can choose the functions $x_1, ..., x_n$ so that they generate the maximal ideal $\mathfrak{m}_p$ of $\mathcal{O}_{X,p}$. We call the set $\{x_i, \partial_i\}_{1\leq i\leq n}$ a local coordinate system *of p on U.*

Proof: [HTT08, Theorem A.5.1]. □

Note that the elements x_i appearing in the local coordinate system above are regular functions $x_i : V \longrightarrow K$, not elements of $\operatorname{End}_K(\mathcal{O}_X(V))$.

It follows from Theorem 1.22 that for any affine open $U \subseteq X$, the ring of differential operators of $\mathcal{O}_X(U)$ is generated by $\mathcal{O}_X(U)$ and $\Theta_X(U)$. This justifies the following definition:

Definition 1.27. Let X be a smooth variety over a field K of characteristic zero. We define the sheaf D_X of *differential operators on* X to be the K-subalgebra of $\mathcal{E}nd_K(\mathcal{O}_X)$ generated by $\mathcal{O}_X$ and Θ_X.

For any point $p \in X$, we may find an affine open $U \subseteq X$ containing p and a local coordinate system $\{x_i, \partial_i\}_{1\leq i\leq n}$ such that

$$D_U = D_X|_U = \bigoplus_{\alpha\in\mathbb{N}^n} \mathcal{O}_X(U)\partial^\alpha$$

by combining Theorems 1.22 and 1.26. When X is not smooth, it is instead necessary to consider the sheaf given locally on open affines by $U \mapsto D_{\mathcal{O}_X(U)}$, where $D_{\mathcal{O}_X(U)}$ is defined as in Definition 1.1. This definition agrees with the one above by the theory we have developed thus far, and as we are only concerned with smooth varieties in these notes, we will always have access to a system of local coordinates.

Alternatively, one can define D_X by gluing over open affines. One does this by setting $\Gamma(U, D_X) = D_{\Gamma(U,\mathcal{O}_X)}$ for each open affine $U \subseteq X$. For this to work, we need the following compatibility result:

Proposition 1.28. Let R be a finitely generated regular K-algebra of dimension n. For nonzero $f \in R$, denote by R_f the localization of R at the set $\{1, f, f^2, ...\}$. Then

$$D_{R_f} \cong R_f \otimes_R D_R \quad \text{and} \quad D^i_{R_f} \cong R_f \otimes_R D^i_R.$$

Proof: Let $\varphi : R \longrightarrow R_f$ be the canonical map. The prime ideals of R_f correspond to the primes in R which avoid f, hence we have an isomorphism $(R_f)_\mathfrak{p} \cong R_{\varphi^{-1}(\mathfrak{p})}$ for each prime $\mathfrak{p} \subseteq R_f$. The local ring $R_{\varphi^{-1}(\mathfrak{p})}$ is regular, hence R_f is regular.

Set $W_f = \{1, f, f^2, ...\} \subseteq R$ so that $R_f = W_f^{-1}R$. By the isomorphism $\Omega_{W_f^{-1}R} \cong W_f^{-1}\Omega_R$ [Har77, Proposition 2.8.3], we have

$$\operatorname{Der}_K(R_f) \cong \operatorname{Hom}_{R_f}(W_f^{-1}\Omega_R, R_f) \cong W_f^{-1}\operatorname{Hom}_R(\Omega_R, R) \cong W_f^{-1}\operatorname{Der}_K(R) \cong R_f \otimes_R \operatorname{Der}_K(R).$$

The R_f-module is generated by R_f and $\operatorname{Der}_K(R_f)$ by Theorem 1.22, hence the above isomorphism extends to an isomorphism $D^i_{R_f} \cong R_f \otimes_R D^i_R$. □

This means that an operator in D_{R_f} extends to an operator in D_R once we multiply by a large enough power of f.

It is worth noting that $\Gamma(X, D_X)$ generally fails to embed in $\operatorname{End}_K(\mathcal{O}_X(X))$ when X is not affine, which explains why we must define differential operators locally on affine opens. We conclude this section on differential operators with an example demonstrating this failure.

Example 1.29. Let $X = \mathbb{P}^1_K$ and let $U_0 = \mathbb{A}^1_K$ and $U_1 = \mathbb{A}^1_K$ denote the standard affine opens of X. If x_0 is the coordinate on U_0 and x_1 the coordinate on U_1, then $\Gamma(U_0, D_X)$ is the Weyl algebra generated by $\hat{x}_0, \partial_0$ and $\Gamma(U_1, D_X)$ is the Weyl algebra generated by $\hat{x}_1, \partial_1$. We may view the sheaf D_X to be the sheaf obtained by gluing $D_X|_{U_0}$ and $D_X|_{U_1}$ over $U_0 \cap U_1$, and hence a global differential operator $\theta \in \Gamma(X, D_X)$ is fully specified by a pair (θ_0, θ_1) of two elements $\theta_0 \in \Gamma(U_0, D_X)$ and $\theta_1 \in \Gamma(U_1, D_X)$ such that $\theta_0 = \theta_1$ on $U_0 \cap U_1$.

We change coordinates from U_0 to U_1 via $x_0 \mapsto x_1^{-1}$. To express ∂_1 in terms of $\hat{x}_0, \partial_0$ on the open set $U_0 \cap U_1$ we use the chain rule:

$$\partial_1 = \frac{\partial}{\partial x_1} = \frac{\partial}{\partial x_0}\frac{d\hat{x}_0}{dx_1} = -\hat{x}_1^{-2}\partial_0 = -\hat{x}_0^2\partial_0.$$

Two differential operators

$$\theta_0 = \sum_{i=1}^{n} a_i \hat{x}_0^{b_i}\partial_0^{c_i} \quad \text{and} \quad \theta_1 = \sum_{j=1}^{m} \alpha_j \hat{x}_1^{\beta_j}\partial_1^{\gamma_j}$$

are therefore equal on $U_0 \cap U_1$ if and only if

$$\sum_{i=1}^{n} a_i \hat{x}_0^{b_i}\partial_0^{c_i} \;=\; \sum_{j=1}^{m} \alpha_j \hat{x}_0^{-\beta_j}\left(-\hat{x}_0^2\partial_1\right)^{\gamma_j}.$$

Determining whether two such arbitrary operators agree on $U_0 \cap U_1$ is quite difficult in general, as it involves expanding multiple terms of the form $(-\hat{x}_0^2\partial_0)^{\gamma}$ at once. However, we can use this restriction criterion to easily construct an infinite set of K-linearly independent global differential operators. Define $\delta = -\hat{x}_0^2\partial \in \Gamma(U_0, D_X)$. Then δ^n is equal to ∂_1^n on $U_0 \cap U_1$ for any $n \in \mathbb{N}$, and so the set $\{(\delta^n, \partial_1^n)\}$ is a K-linearly independent set of global differential operators. This means $\Gamma(X, D_X)$ is infinite dimensional as a K-vector space.

Since $\mathrm{End}_K(\mathcal{O}_X(X)) = \mathrm{End}_K(K) = K$ is a 1-dimensional K-vector space, there is no embedding $\Gamma(X, D_X) \longrightarrow \mathrm{End}_K(\mathcal{O}_X(X))$.

1.4 A Word Regarding Non-Regular *K*-Algebras

To conclude our discussion of the ring of differential operators, we say a brief word about the singular case. There is still an "R-linear" way to compute the modules D^i_R even when R is not regular. We loosen our assumptions on R and once again take R to be an algebra over another commutative ring A. Taking cues from the characterization of $\mathrm{Der}_K(R)$ in terms of Kähler differentials, we define

$$P^i_{R/A} = \frac{R \otimes_A R}{\Delta^{i+1}_{R/A}} \tag{3}$$

to be the module of ith principal parts of R over A, with $\Delta_{R/A}$ as in Definition 1.10. One can then prove that

$$D^i_{R/A} \cong \mathrm{Hom}_R(P^i_{R/A}, R).$$

Thus the functor $R \mapsto D^i_{R/A}$ is represented by $P^i_{R/A}$. This construction can be found in [Moo04], who attributes it to Grothendieck.

2 *D*-Modules: Basic Definitions and Facts

We start with the definition of a D-module.

Definition 2.1. Let X be a smooth variety over a field K. A left (or right) D-module over X, or a D_X-module, is a quasi-coherent $\mathcal{O}_X$-module $\mathcal{M}$ together with a left (or right) action by D_X. We say that $\mathcal{M}$ is a *coherent* D_X-module if it is locally finitely generated over D_X.

In the affine case, a D-module corresponds to a module M over a ring of differential operators, i.e. a left or right D_R-module, via the correspondence between $\Gamma(X, \mathcal{O}_X)$-modules and quasi-coherent sheaves of $\mathcal{O}_X$-modules (see Example 2.9). When working over an affine variety $\operatorname{Spec} R = X$, it therefore suffices (and is more convenient) to study $M = \Gamma(X, \mathcal{M})$ rather than $\mathcal{M}$ itself.

Note that a coherent D_X-module is *not* necessarily coherent as an $\mathcal{O}_X$-module. For instance, $A_n(K) = \Gamma(X, D_X)$ is the nth Weyl algebra when $X = \operatorname{Spec} K[x_1, ..., x_n]$, and though $A_n(K)$ is trivially finitely generated as a module over itself, it is certainly not finitely generated as a $\Gamma(X, \mathcal{O}_X) = K[x_1, ..., x_n]$-module as there is no way to increase the degree of an operator via the action of a polynomial.

We start this section with several examples before discussing the basic theory relating to the structure of D-modules.

2.1 Examples of *D*-modules

Let R be a regular finitely generated K-algebra. We start with a trivial example.

Example 2.2. Every ring is a module over itself, so D_R is a left D_R-module as are all of its left ideals. The polynomial ring R is also a left D_R-module, where the left action of an operator $\alpha \in D_R$ on $f \in R$ is given by applying α to f, i.e. $\alpha \cdot f = \alpha(f)$.

This is quite unremarkable, so we quickly move on to some more interesting examples.

Example 2.3. Let $R = K[x]$, so that $D_R = A$ is the first Weyl algebra. Suppose $I = A\partial$ and $J = A\hat{x}$ be the left ideals of A generated by ∂ and $\hat{x}$ respectively and let $M = A/I$ and $N = A/J$. These are quotients of left A-modules and are therefore themselves left A-modules. As K-vector spaces, it is clear that $M \cong K[\hat{x}]$ and $M \cong K[\partial]$.

To understand the A-action on M, it suffices to understand the action of $\hat{x}$ and ∂ on the basis $\{1 + I, \hat{x} + I, \hat{x}^2 + I, ...\}$ of M. The operator $\hat{x}$ acts by multiplication; it's an infinite Jordan block with one's along the upper diagonal and zeros elsewhere. Let us now consider the action of ∂ on M. Since $\partial\hat{x} = 1 + \hat{x}\partial$ and $\hat{x}\partial \in I$, we have that $\partial(\hat{x} + I) = 1 + I$. Similarly, $\partial(\hat{x}^k + I) = \partial(\hat{x}^k) + I = k\hat{x}^{k-1}$. Thus, the map K-linear map $K[x] \longrightarrow M$ given by $x \mapsto \hat{x}$ is compatible with the action of A, and is hence an isomorphism of left A-modules. The K-linear isomorphism $N \longrightarrow K[\partial]$ can be used to identify the action of A on N with $K[\partial]$; ∂ acts on $K[\partial]$ by multiplication and $\hat{x} \cdot \partial = -1$. This also gives us an isomorphism of left A-modules.

In addition to their structural similarities, the modules M and N are related in a slightly deeper way. Suppose we are given a ring S, an automorphism $\sigma : S \longrightarrow S$ and a left S-module P. The *twist* of P by σ is denoted P_σ. It is isomorphic to P as an Abelian group and has a left S-action defined by $s \cdot p = \sigma(s)p$ for $p \in P$ and $s \in S$.

Consider now the automorphism F of A defined $F(x) = -\partial$ and $F(\partial) = x$. The kernel of F composed with the projection $A \longrightarrow A/I$ is $F^{-1}(I) = A\hat{x} = J$, so

$$A/F^{-1}(I) \cong M$$

as A-modules by the first isomorphism theorem. Furthermore, the projection $\pi : A \longrightarrow (A/I)_F$ satisfies

$$\pi(a) = a \cdot \pi(1) = F(a) + I$$

for $a \in A$ and has kernel $F^{-1}(J)$. We therefore have isomorphisms

$$N_F = (A/I)_F \cong A/F^{-1}(I) \cong M.$$

This construction is important enough that we give it a name.

Definition 2.4. Let M be a left A_n-module. The *Fourier transform* $\widehat{M}$ of M is defined as follows. As an additive group, $\widehat{M} = M$, while the actions of $\hat{x}_i$ and ∂_{x_i} on $u \in M$ are given by

$$\hat{x}_i \cdot u = -\partial_{x_i}(u), \qquad \text{and} \qquad \partial_{x_i} \cdot u = \hat{x}_i(u).$$

This is equivalent to defining $\widehat{M} = M_F$ with F defined in Example 2.3.

This will appear again in the final section.

Example 2.5. Let $K = \mathbb{C}$, denote by A the Weyl algebra over $\mathbb{C}$, and fix a subset $U \subseteq \mathbb{C}$ open with respect to the Euclidean topology. Every holomorphic function on $\mathbb{C}$ is analytic, and therefore the set $\mathcal{H}(U)$ of holomorphic functions on U is a left A-module. Somewhat more surprising is the fact that it is not a torsion module; one can show that the function $h(x) = \exp(\exp(z))$ is not killed by any element of A for instance [Cou95, Proposition 5.3.2].

Example 2.6 (Module Associated to a Differential Equation)**.** Let $K = \mathbb{R}$, denote by A_n the nth Weyl algebra and fix a set $U \subseteq \mathbb{R}^n$. The set $\mathcal{C}^\infty(U)$ of infinitely differentiable functions in $x_1, ..., x_n$ is then an A_n module.

Consider now an arbitrary operator $P = \sum_{i=1}^m g_{\alpha_i} \partial^{\alpha_i} \in A_n$ where $\alpha_i \in \mathbb{N}^n$ is a multi-index for each $1 \leq i \leq n$. This operator gives us a differential equation:

$$P(f) = \sum_{i=1}^m g_{\alpha_i} \partial^{\alpha_i}(f) = 0$$

where $f \in C^\infty(U)$. We can similarly define a system of differential equations

$$P_1(f) = ... = P_k(f) = 0 \tag{4}$$

given $P_1, ..., P_k \in A_n$. The $\mathbb{R}$-vector space of solutions to this system is certainly not an A_n-module, if f satisfies the system there is no expectation that $\partial_{x_i}(f)$ does as well for instance, but it does nonetheless admit a nice description via the theory of A_n-modules.

Let $J = \sum_{i=1}^k A_n P_k$ be the left ideal generated by $P_1, ..., P_k$ and set $M = A_n/J$. We say that M is the A_n-module associated to the system (4). We will show that the set of polynomial solutions to (4) is isomorphic to $\mathrm{Hom}_{A_n}(M, \mathbb{R}[x_1, ..., x_n])$ as a $\mathbb{R}$-vector space.

First, consider a polynomial solution $f \in \mathbb{R}[x_1, ..., x_n]$ to (4), and associate to f the A_n-module homomorphism $\varphi_f : A_n \longrightarrow \mathbb{R}[x_1, ..., x_n]$ defined by $1 \mapsto f$. If $Q \in J$, then $Q(f) = 0$, so $\varphi_f(Q) = 0$ and hence φ_f induces a map $\overline{\varphi_f} : M \longrightarrow \mathbb{R}[x_1, ..., x_n]$.

Consider now the $\mathbb{R}$-linear map $f \mapsto \overline{\varphi_f}$ taking a polynomial solution of (4) to its associated A_n-module homomorphism. The map $\sigma \mapsto \sigma(1)$ which sends a homomorphism $\sigma : M \longrightarrow \mathbb{R}[x_1, ..., x_n]$ to its evaluation at $1 \in M$ is the inverse of $f \mapsto \overline{\varphi_f}$, hence it is an isomorphism.

These examples have all been of left A_n-modules, but we can turn left modules into right modules and vice versa. Let R be a regular K-algebra of finite type as in Theorem 1.22.

Example 2.7. (Swapping Left and Right Modules) Consider an operator $P \in D_R$ given by $P = \sum_\alpha \hat{f}_\alpha \partial^\alpha$. It's *formal adjoint* is the operator

$$^tP := \sum_\alpha (-\partial)^\alpha \hat{f}_\alpha \in D_R.$$

This satisfies $^t(PQ) =^t Q\, ^tP$, so $P \mapsto {}^tP$ is an anti-automorphism of D_R. Given a left D_R-module M, we can obtain a right D_R-module ^tM which is isomorphic to M as an Abelian group and whose D_R-action is given by $u \cdot P = {}^tPu$. We can do something similar to obtain a left module from a right module.

This notion depends on choice of local coordinates, and therefore does not extend to non-affines D_X-modules. The correct globalization of this process involves the canonical sheaf, see [HTT08, Chapter 1.2].

Now let X be a smooth variety over K.

Example 2.8. A necessary and sufficient condition for a sheaf $\mathcal{F}$ of $\mathcal{O}_X$-modules to be affine is that for any affine $U \subseteq X$, $\mathcal{F}|_U \cong \tilde{M}$ where $M = \Gamma(U, \mathcal{F})$ (see [Har77, Chapter 2.5]). If $U \subseteq X$ is affine and $f \in \mathcal{O}_X(U)$, then

$$\Gamma(D(f), D_X|_U) = D_{\mathcal{O}_X(U)_f} \cong \mathcal{O}_X(U)_f \otimes_{\mathcal{O}_X(U)} D_{\mathcal{O}_X(U)}$$

by Proposition 1.28, so $D_X|_U \cong \Gamma\widetilde{(U, D_X)}$. This implies that D_X is itself a left D-module, and we can similarly see that $\mathcal{O}_X$ is a left D_X-module. Indeed, for any open affine U, the algebra $D_X(U)$ acts on $\Gamma(U, D_X)$ and $\Gamma(U, \mathcal{O}_X)$ by the construction in Section 1.3.2.

Example 2.9. When $X = \operatorname{Spec} R$ is affine, every left D_X-module $\mathcal{M}$ corresponds to a left D_R-module via $\mathcal{M} \mapsto \Gamma(X, \mathcal{M})$. Examples 2.3 through 2.7 are therefore all examples of D-modules over $\mathbb{A}^n_K$ once we pass to the associated sheaf.

2.2 Filtrations

Though we have neither commutativity nor a graded structure, we can extend the tools of used in commutative algebra to study graded modules to better understand the structure of D-modules. This is accomplished by associating a graded commutative ring to D_R and a compatible graded module to a D_R-module M through the use of filtrations. These methods will be especially fruitful in the study of modules over the Weyl algebra.

This is a brief overview of some definitions concerning filtered K-algebras, tailored to the purposes of this essay. We are primarily interested in good filtrations of finitely generated A_n-modules, as these provide us with sufficient conditions to discuss dimension. A more general treatment suitable to the case of $\mathcal{D}_X$-modules over a scheme X can be found in Chapter 1 of [Gin98], which largely serves as the inspiration for this section. Though all of our statements deal with left modules, everything holds if we replace “left” with “right” and make the obvious, necessary changes.

Definition 2.10. Let R be a K-algebra. We say R is a *filtered K-algebra* if it comes equipped with a collection $\{F_i\}_{i\in\mathbb{N}}$ of K-vector spaces such that

(1) $K = F_0 \subset F_1 \subset F_2 \subset ... \subset R$

(2) $F_i \cdot F_j \subseteq F_{i+j}$.

(3) $R = \bigcup_{i\geq 0} F_i$, (we say the filtration is *exhausting*)

When equipped with a filtration, R is said to be a *filtered K-algebra.* We often write this as a pair $(R, F_\bullet)$, set $F_{-1} = \{0\}$ and iterate over $\mathbb{Z}$ rather than $\mathbb{N}$.

Remark 2.11 (Definition Cont.)**.** Let $(R, F_\bullet)$ be as in the above definition. The collection of sets $\{F^i + r\}_{i\in\mathbb{Z}, r\in R}$ form the basis of a topology on R. With this in mind, it is often convenient to impose two additional conditions:

(4) $\bigcap_{i\geq -1} F_i = \{0\}$, which is equivalent to say that the topology induced by $F_\bullet$ is separating,

(5) R is complete with respect to this topology.

We also have a notion of a filtered ring in which we replace the K-vector spaces with abelian groups, but in this essay we will only be concerned with filtered K-algebras.

Example 2.12. The collection $D_R^\bullet = \left\{D_R^k\right\}_{k\in\mathbb{N}}$ is a filtration of D_R. Requirement (1) holds by Lemma 1.4, requirement (2) by Proposition 1.7 (be) and requirement (3) by definition of D_R. This is called the *order filtration* on D_R. Note that $D_R^0 = R$, and is therefore an infinite dimensional K-vector space.

Example 2.13. The n^{th} Weyl algebra A_n is comes equipped with another filtration, the *Bernstein filtration.* We denote this filtration $\mathcal{B} = \{B_k\}_{k\geq 0}$ where $B_k = \{D \in A_n \mid \deg(D) \leq k\}$. The operators $\hat{x}^\alpha \partial^\beta$ of degree at most k form a basis for B_k over K, and hence the Bernstein filtration has the added benefit that each component is finite dimensional.

Example 2.14. Suppose $R = \bigoplus_{i\in\mathbb{N}} R_i$ is a graded ring. Then $(R, F_\bullet)$ is a filtered K-algebra with respect to the filtration $F_k = \bigoplus_{i=0}^k R_i$.

Definition 2.15. Let $(R, F_\bullet)$ be a filtered K-algebra. The *associated graded K-algebra*, $\operatorname{gr}^{F_\bullet} R$, is defined

$$\operatorname{gr}^{F_\bullet} R = \bigoplus_{i=0}^{\infty} F_i / F_{i-1}.$$

When the filtration is known, we write $\operatorname{gr} R$. For any $r \in F_i$, we denote by $\sigma_i(r)$ its image in F_i/F_{i-1} and say $\sigma_i(r)$ is the i^{th} *principal symbol* of r. The associated graded ring to the filtration given in Example 2.14 recovers the original graded ring, as one might hope.

We use the principal symbol maps σ_i to define an algebra structure on $\operatorname{gr}^{F_\bullet} R$. A *homogeneous element* of $\operatorname{gr}^{F_\bullet} R$ is any operator $d \in \operatorname{gr}^{F_\bullet} R$ such that $d = \sigma_k(a)$ for some $a \in F_k$. Given two homogeneous elements $\sigma_i(a)$ and $\sigma_j(b)$, we define their product by

$$\sigma_i(a) \cdot \sigma_j(b) = \sigma_{i+j}(a \cdot b).$$

Extending this multiplication to all of $\operatorname{gr}^{F_\bullet} R$ by distributivity makes $\operatorname{gr}^{F_\bullet} R$ into a graded K-algebra whose homogeneous components are the individual summands F_k/F_{k-1}.

Example 2.16. Let $S_n = \operatorname{gr}^{\mathcal{B}} A_n$. Then the graded algebra S_n is isomorphic to $K[y_1, ..., y_{2n}]$.

The conceptual sketch of this statement is perhaps more enlightening than the full proof. Since we have surjective maps $\pi_k : A_n \longrightarrow B_k \xrightarrow{\sigma_k} B_k/B_{k-1}$, S_n is generated as an algebra by the images of elements $x_1, ..., x_n, \partial_1, ..., \partial_n \in A_n$. The only thing preventing us from defining a isomorphism $K[y_1, ..., y_{2n}] \longrightarrow S_n$ sending $y_i \mapsto x_i$ and $y_{i+n} \mapsto \partial_i$ for $1 \leq i \leq n$ is a possible lack commutativity, however, we saw in Corollary 1.8 that the graded ring associated to the order filtration is commutative. This allows us to define a surjective homomorphism $K[y_1, ..., y_{2n}] \longrightarrow S_n$. Since there are no additional relations between the generators $x_1, ..., x_n, \partial_1, ..., \partial_n$, the kernel of this map is trivial and hence we have an isomorphism.

Definition 2.17. Let $(R, F_\bullet)$ be a filtered K-algebra and M a left R-module. A *filtration of* M compatible with $F_\bullet$is a family $\Gamma = \{\Gamma_0\}_{i \geq 0}$ of K-vector spaces satisfying

(1) $\Gamma_0 \subseteq \Gamma_1 \subseteq \Gamma_2 \subseteq ... \subseteq M$,

(2) $F_i\Gamma_j \subseteq \Gamma_{i+j}$.

(3) $M = \bigcup_{i \geq 0} \Gamma_i$

Such a module is said to be *filtered*, and as with algebras, we set $\Gamma_{-1} = 0$. The *associated graded module* to M is

$$\operatorname{gr}^{\Gamma} M = \bigoplus_{i=0}^{\infty} \Gamma_i/\Gamma_{i-1}$$

and is a graded gr R module.

The associated grading can tell us something about its filtered module.

Theorem 2.18. *Suppose that R is a filtered K-algebra with filtration $F_\bullet$ such that $S = \operatorname{gr}^{F_\bullet} R$ is Noetherian. Let M be a left R-module with filtration $\Gamma = \{\Gamma_i\}_{i \geq 0}$. If $\operatorname{gr}^{\Gamma} M$ is a Noetherian then so is M.*

Proof: Let $N \subseteq M$ be a R-submodule of M. We prove that it is finitely generated. Define $\Gamma'_i = N \cap \Gamma_i$ for $i \geq 0$. The collection $\Gamma' = \{\Gamma'_i\}$ is then a filtration of N, which we call the *induced filtration of* N *by* Γ. The inclusions $\Gamma'_i \subseteq \Gamma_i$ give us an inclusion $\operatorname{gr}^{\Gamma'} N \subseteq \operatorname{gr}^{\Gamma} M$, and since $\operatorname{gr}^{\Gamma} M$ is Noetherian, $\operatorname{gr}^{\Gamma'} N$ must be a finitely generated as an S-module.

Let $\{c_1, ..., c_r\}$ be a generating set for $\operatorname{gr}^{\Gamma'} M$. We assume that each c_i is homogeneous without loss of generality; each c_i is a linear sum of finitely many homogeneous elements and we can therefore replace each c_i by its homogeneous components without compromising the finiteness of our generating set. For each c_i we can therefore find some integer k_i and some $u_i \in \Gamma'_{k_i}$ such that $\mu_{k_i}(u_i) = c_i$. Let $m = \max\{k_1, ..., k_r\}$, and note that $u_i \in \Gamma'_m$ for each $1 \leq i \leq r$. We show that Γ'_m generates N.

Suppose $v \in \Gamma_\ell$. If $\ell \leq m$ then $v \in \Gamma'_\ell \subseteq \Gamma'_m$, and hence v is in the R-submodule of M generated by Γ'_m. Suppose now that $\ell > m$ and $\Gamma_{\ell-1}$ is contained in the R-linear span of Γ'_m. Because $\{\mu_{k_1}(u_1), ..., \mu_{k_r}(u_r)\}$ generates $\operatorname{gr}^{\Gamma'} N$ as an S-module, there exist $a_1, ..., a_r$ such that

$$\mu_\ell(v) = \sum_{i=1}^{r} \sigma_{\ell-k_i}(a_i)\mu_{k_i}(u_i).$$

Hence

$$\mu_\ell\left(v - \sum_{i=1}^{r} a_i u_i\right) = 0$$

$$v' = v - \sum_{i=1}^{r} a_i u_i \in \Gamma'_{\ell-1}.$$

The element v is a linear sum of elements in Γ'_m if and only if v' is too. However, $v' \in \Gamma'_{\ell-1}$ and is therefore in the R-linear span of Γ'_m by the inductive hypothesis. Hence $v \in R \cdot \Gamma'_m$, and since every element of N is contained in Γ'_ℓ, Γ'_m generates N.

It is left to show that there is a finite subset of Γ'_m which generates N. However, Γ'_m is a finite dimensional K-vector space. Any K-basis for Γ'_m will generate all of Γ'_m and will therefore serve as a set of generators for N. □

Note that the set $\{u_1, ..., u_r\}$ in the above proof is not necessarily a generating set for N. The induction step gives us an algorithm for writing any $v \in \Gamma'_\ell$ in terms of the u_i only in the case that $\ell > \max\{\deg(u_1), ..., \deg(u_r)\}$.

We have the following immediate corollary.

Corollary 2.19. The nth Weyl algebra A_n is left Noetherian.

Proof: The associated graded ring of A_n with respect to the Bernstein filtration is the polynomial ring in two variables by Example 2.16, which is Noetherian. □

As mentioned in the introduction to this section, these statements hold if we replace "left" by "right" and make the necessary adjustments, meaning A_n is also right Noetherian. This is quite convenient, for it means any finitely generated left or right A_n-module is automatically Noetherian.

The converse of Theorem 2.18 need not always hold, that is, it need not be the case that $\operatorname{gr}^\Gamma M$ is finitely generated even if M is finitely generated. We therefore distinguish filtrations which produce finitely generated associated graded modules.

Definition 2.20. Let M be a left module over a filtered K-algebra $(R, F_\bullet)$. A filtration Γ of M is said to a *good filtration* with respect to $F_\bullet$ if $\operatorname{gr}^\Gamma M$ is finitely generated. Good filtrations provide a framework to discuss the dimension of modules over the Weyl algebra.

Good filtrations always exist for finitely generated modules.

Proposition 2.21. Let $(R, F_\bullet)$ be a filtered K-algebra and M be a finitely generated left R-module. Then there exists a good filtration Γ of M compatible with $F_\bullet$.

Proof: Let $u_1, ..., u_r$ be a generating set for M over R and define $\Gamma_k = \sum_{i=0}^{r} F_k u_i$. Then $\operatorname{gr}^\Gamma M$ is finitely generated over $\operatorname{gr}^{F_\bullet} R$ by the images of $u_1, ..., u_k$ in Γ_k. □

We end the section on filtrations by stating two propositions, both of which are included primarily for convenient use in the discussion of holonomic modules over A_n. One provides a criterion for easily checking whether a filtration is good, and the other allows us to compare two good filtrations of a module.

Proposition 2.22. Let M be a left module over a filtered K-algebra $R, F_\bullet)$. A filtration Γ of M with respect to $F_\bullet$ is good if and only if there exists an integer k_0 such that $\Gamma_{i+k} = F_i\Gamma_k$ for all $k \geq k_0$.

Proof: [Cou95, Proposition 8.3.1] □

This criterion is useful for determining both good and bad filtrations.

Example 2.23. Consider the Bernstein filtration $\mathcal{B} = \{B_i\}_{i\in\mathbb{N}}$ on the Weyl algebra A_n. Set $\Gamma_i = B_{2i}$. We then have that $B_i\Gamma_k = B_{i+2k} \neq B_{2(i+k)} = \Gamma_{i+k}$, so Γ_i is not a good filtration of A_n with respect to the Bernstein filtration.

Proposition 2.24. Let M be a left module of the filtered K-algebra $(R, F_\bullet)$. Suppose that Γ and Λ are two filtrations of M with respect to $F_\bullet$. The following statements are true.

(a) If Γ is good with respect to $F_\bullet$ then there exists some k_0 such that $\Lambda_i \subseteq \Lambda_{i+k_0}$ for all $i \in \mathbb{N}$.

(b) If both Γ and Λ are good with respect to $F_\bullet$, then there exists some k_1 such that $\Lambda_{i-k_1} \subseteq \Gamma_i \subseteq \Lambda_{i+k_1}$.

Proof: [Cou95, Proposition 8.3.2] □

2.3 Modules over the Weyl algebra

The finite dimensionality of each component in the Bernstein filtration makes it distinctly nice because of the following property. For any module M over the Weyl algebra with a good filtration Γ, the K-vector spaces Γ_i are each finite dimensional. This means we can apply results regarding Hilbert polynomials to modules over the Weyl algebra, and thereby obtain a robust dimension theory with relatively little headache.

Throughout this section K is a field of characteristic 0, $A_n = D_{K[x_1,...,x_n]}$ is the nth Weyl algebra, S_n is the associated graded ring to A_n with respect to the Bernstein filtration $\mathcal{B}$, and M is a finitely generated left A_n module. We work almost entirely with left A_n-modules in this section, but all results in this section hold if "left" is replaced with "right" and the obvious modifications are made.

2.3.1 Dimension

The primary goal of this section is a proof of Bernstein's Inequality, a striking example of how the theory of D-modules can drastically differ from that of modules over commutative rings. To accomplish this, it is necessary to discuss several basic facts regarding the dimension of modules over the Weyl algebra, theory which relies on dimension theory from commutative algebra. We brazenly omit proofs and discussion of these facts in eternal deference to Atiyah-Macdonald [AM16].

Recall that if $M = \oplus_{i\geq 0} M_i$ is a finitely generated graded module over a polynomial ring $K[x_1, ..., x_m]$, then there exists a polynomial $\chi(t) \in \mathbb{Q}[t]$ and a positive integer N such that

$$\sum_{i=0}^{t} \dim_K(M_i) = \chi(t)$$

for all $t \geq N$. We typically suppress N from our notation and simply write "for all $t \gg 0$" to mean "for all t sufficiently large". The polynomial $\chi(t)$ is called the Hilbert polynomial of M.

If M is a finitely generated left A_n-module then there exists a filtration Γ of M which is good with respect to the Bernstein filtration by Proposition 2.21. The associated graded module $\operatorname{gr}^\Gamma M$ is then seen to be Noetherian since it is finitely generated over $S_n = \operatorname{gr}^\Gamma A_n$, a Noetherian ring. This means the Hilbert polynomial for $\operatorname{gr}^\Gamma M$ exists, and we denote it by $\chi(t, \Gamma, M) \in \mathbb{Q}[t]$. This discussion leads us to the following definition.

Definition 2.25. Let M be a finitely generated left A_n-module equipped with a good filtration Γ with respect to the Bernstein filtration. Denote by $\chi(t, \Gamma, M)$ the Hilbert polynomial of $\operatorname{gr}^\Gamma M$. Let a be the leading coefficient of $\chi(t, \Gamma, M)$ and let d be its degree. The *dimension* $d(M)$ of M is d and the *multiplicity* $m(M)$ of M is $d! \cdot a$. Both of these are nonnegative integers.

See [AM16] for details, or [Cou95, Chapter 9] for a discussion tailored specifically to modules over the Weyl algebra. The latter sources also provides a brief argument demonstrating that the definitions of dimension and multiplicity do not depend on the choice of good filtration.

Example 2.26. It is well known that the Hilbert polynomial of the polynomial rink $K[x_1, ..., x_m]$ is degree m. Hence the Hilbert polynomial of $S_n = K[y_1, ..., y_{2n}]$ is degree $2n$ and $d(A_n) = 2n$. By this same argument, $d(K[x_1, ..., x_n]) = n$.

Proposition 2.27. Let M be a finitely-generated left A_n-module and $N \subseteq M$ a submodule. Then

(a) $\dim(M) = \max\{d(N), d(M/N)\}$

(b) We have

$$m(M) = \begin{cases} m(N) + m(M/N) & \text{if } d(N) = d(M/N) \\ m(N) & \text{if } d(N) > d(M/N) \\ m(M/N) & \text{if } d(N) < d(M/N) \end{cases}.$$

Proof:

(a) Let us first see how the Hilbert polynomials of M, N and M/N related. Denote by S_n the associated graded ring of A_n, and let Γ be a good filtration of M with respect to $\mathcal{B}$. Let Γ' and Γ'' be the induced filtrations for N and M/N. We then obtain the following short exact sequence of associated graded S_n-modules:

$$0 \longrightarrow \operatorname{gr}^{\Gamma'} N \longrightarrow \operatorname{gr}^\Gamma M \longrightarrow \operatorname{gr}^{\Gamma''} M/N \longrightarrow 0.$$

We know $\operatorname{gr}^\Gamma M$ is a finitely generated S_n-module since Γ is good, hence $\operatorname{gr}^{\Gamma''} M/N$ is also finitely generated since it is isomorphic to a quotient of $\operatorname{gr}^\Gamma M$. Likewise, since S_n is Noetherian and $\operatorname{gr}^{\Gamma'} N$ is isomorphic to a submodule of $\operatorname{gr}^\Gamma M$, $\operatorname{gr}^{\Gamma'} N$ is finitely generated. This tells us that Γ' and Γ'' are both good filtrations.

Now consider the short exact sequence of vector spaces

$$0 \longrightarrow \Gamma'_k/\Gamma'_{k-1} \longrightarrow \Gamma_k/\Gamma_{k-1} \longrightarrow \Gamma''_k/\Gamma''_{k-1} \longrightarrow 0$$

for $0 \leq k$. By the rank-nullity theorem, $\dim_K \Gamma_k/\Gamma_{k-1} = \dim_K \Gamma'_k/\Gamma'_{k-1} + \dim_K \Gamma''_k/\Gamma''_{k-1}$, so

$$\sum_{k=0}^{\infty} (\dim_K \Gamma_k/\Gamma_{k-1}) = \sum_{k=0}^{\infty} \left(\dim_K \Gamma'_k/\Gamma'_{k-1} + \dim_K \Gamma''_k/\Gamma''_{k-1}\right)$$

and thus for $s \gg 0$ we get

$$\chi(s, \Gamma, M) = \chi(s, \Gamma', N) + \chi(s, \Gamma'', M/N).$$

As all of the above are polynomials with positive leading coefficients,

$$\deg\big(\chi(s, \Gamma', N) + \chi(s, \Gamma'', M/N)\big) = \deg\big(\chi(s, \Gamma', N)\big) + \big(\chi(s, \Gamma'', M/N)\big)$$

and hence

$$d(M) = \max\{d(N), d(M)\}.$$

(b) If $d(N) > d(M)$, then the degree of $\chi(s, \Gamma', N)$ is strictly larger than the degree of $\chi(s, \Gamma'', M/N)$, and hence the leading coefficient of $\chi(s, \Gamma, M)$ is equal to the leading coefficient of $\chi(s, \Gamma', N)$ by what we proved in (a). The same argument works for the case that $d(N) < d(M)$. If instead $d(M/N) = d(N)$, then the polynomials $\chi(s, \Gamma, M), \chi(s, \Gamma', M)$ and $\chi(s, \Gamma'', N)$ all have the same degree. This then implies that the leading term of $\chi(s, \Gamma, M)$ is equal to the sum of the leading terms of $\chi(s, \Gamma', N)$ and $\chi(s, \Gamma'', M/N)$.

□

Corollary 2.28. Let M be a finitely generated A_n-module. Then $d(M) \leq 2n$.

Proof: Let $\{u_1, ..., u_r\}$ be a generating set over A_n for M. There then exists a surjective homomorphism $\phi : A_n^{\oplus r} \longrightarrow M$. Proposition 2.27 then tells us that $d(A_n^{\oplus r}) = \max\{d(M), d(\ker\phi)\}$.

We claim that $d(A_n^{\oplus r}) = 2n$. Indeed, we have seen that $d(A_n) = 2n$, and there exists an exact sequence

$$0 \longrightarrow A_n \longrightarrow A_n^{\oplus r} \longrightarrow A_n^{\oplus(r-1)} \longrightarrow 0$$

from which we get that $(A_n^{\oplus r}) = \max\{d(A_n), d(A_n^{\oplus(r-1)})\}$. Induction on r then gives us the desired result, hence $\max\{d(M), d(\ker\phi)\} = 2n$. We conclude $d(M) \leq 2n$. □

2.3.2 Bernstein's Inequality

We now prove Bernstein's inequality, which serves as a striking example of the difference between the structure of modules over the polynomial ring and modules over the Weyl algebra.

Theorem 2.29 (Bernstein's Inequality). *If M is a finitely-generated left $A_n(K)$-module, then either $n \leq \dim(M)$ or $M = 0$.*

Proof: Let $\mathcal{B} = \{B_k\}_{k\geq 0}$ be the Bernstein filtration. Fix a generating set $u_1, ..., u_r$ for M over A_n and let Γ be the good filtration obtained by setting $\Gamma_k = \sum_{i=1}^{r} B_k u_i$, as in the proof of Proposition 2.21. Finally, let $\chi(t) = \chi(t, \Gamma, M)$ be the Hilbert polynomial of M.

We first show that the K-vector space B_i embeds in $\mathrm{Hom}_K(\Gamma_i, \Gamma_{2i})$ for each $i \geq 0$. For $a \in B_i$ define $\phi_a : \Gamma_i \longrightarrow \Gamma_{2i}$ by $u \mapsto au$ and let $\phi : B_i \longrightarrow \mathrm{Hom}_K(\Gamma_i, \Gamma_{2i})$ be the K-linear map $a \mapsto \phi_a$, noting that ϕ is injective exactly when $a\Gamma_i \neq 0$ for any $0 \neq a \in B_i$. We prove that $a\Gamma_i$ is never 0 by induction on i.

For $i = 0$ we have $B_0 = K$, and hence ϕ is injective exactly when $\Gamma_0 \neq 0$. Since $u_1, ..., u_r \in \Gamma_0$, this is satisfied.

Assume now that ϕ is injective for all $1 \leq j < i$, that is, if $0 \neq b \in B_j$ then $b\Gamma_j \neq 0$. Fix some nonzero $a \in B_i$. The canonical form of a must then include a nonzero term which is a product of either $\hat{x}_\ell$ or ∂_{x_ℓ} for some $1 \leq \ell \leq n$. In particular,

$$[a, P] \neq 0, \quad \text{for some } \ P \in \{\hat{x}_1, ..., \hat{x}_n, \partial_{x_1}, ..., \partial_{x_n}\}.$$

Suppose for the sake of contradiction that $a\Gamma_i = 0$. Since $\deg(P) = 1$, $P\Gamma_{i-1} \subseteq \Gamma_i$, so $a(P\Gamma_{i-1}) = 0$. We then have that

$$[a, P]\Gamma_{i-1} = a(P\Gamma_{i-1}) - P(a\Gamma_{i-1}) = 0. \tag{$*$}$$

However, $\deg([a, P]) \leq \deg(a) - 1$ by Proposition 1.19 (c), so $[a, P]$ is a nonzero element of B_{i-1}. Hence $(*)$ contradicts the inductive hypothesis and $a\Gamma_i \neq 0$. This proves that ϕ is injective for all values $i \geq 0$.

We now prove that $d(M) \geq n$. The injectivity of ϕ implies

$$\dim_K(B_i) \leq \dim_K(\mathrm{Hom}_K(\Gamma_i, \Gamma_{2i})) \tag{$\dagger$}$$

for all $i \geq 0$. Let's examine the RHS of this inequality. It is a fact of elementary linear algebra that $\dim_K(\mathrm{Hom}_K(\Gamma_i, \Gamma_{2i}) = \dim_K(\Gamma_i)\dim_K(\Gamma_{2i})$, hence for $i \gg 0$, $\dim_K(\mathrm{Hom}_K(\Gamma_i, \Gamma_{2i}) = \chi(i)\chi(2i)$.

Now consider the LHS. By definition, the set of all elements of the form $\hat{x}^\alpha\partial^\beta$ with $\alpha, \beta \in \mathbb{N}^n$ satisfying $|\alpha| + |\beta| \leq 2i$ forms a basis for B_i as a K-vector space. A combinatorial argument shows that the number of monomials in k variables of degree at least d is $\binom{k+d}{k}$. Hence $\dim_K(B_i) = \binom{i+2n}{2n}$. Expanding, we see that

$$\binom{i+2n}{2n} = \frac{(i+2n)!}{i!(2n)!} = \frac{1}{(2n)!}(i+2n)(1+2n-1)...(1+2n-(2n-1))$$

is a polynomial in i of degree $2n$. In order for the inequality ($\dagger$) to hold for all values of i, $\chi(i)\chi(2i)$ must likewise be at least degree $2n$. However, $\deg(\chi(i)\chi(2i)) = 2\deg(\chi(i)) = 2d(M)$. This means $2d(M) \geq 2n$, or $d(M) \geq n$ as desired. □

2.3.3 Holonomic Modules

The Bernstein inequality tells us that a nonzero finitely-generated left $A_n(K)$-module M must have dimension at least n. Those modules of minimal dimension are called *holonomic modules*. Holonomic modules turn out to have particularly nice properties; for instance, they are preserved under inverse and direct images, as we shall see in a later section.

Definition 2.30. A finitely generated left $A_n(K)$-module M is said to be *holonomic* if either $M = 0$ or $\dim(M) = n$.

Examples are easy to identify thanks to Bernstein. We know that $R = K[x_1, ..., x_n]$ is holonomic since $\dim K[x_1, ..., x_n] = n$, and furthermore, both I and R/I are holonomic when I is any proper ideal of R by Proposition 2.27. As another example, in the case that $n = 1$, for any nonzero ideal $I \subseteq A_1$ we have that $\dim(A_1/I) \leq 1$ by Proposition 2.27. We know A_n/I is nonzero since I is proper, hence $\dim(A_1/I) = 1$ by Bernstein's inequality.

Proposition 2.31. The following are true.

(a) Submodules and quotients of holonomic A_n-modules are holonomic.

(b) Direct sums of holonomic A_n-modules are holonomic.

Proof: Statement (a) follows from Bernstein's inequality and the fact that for any finitely generated A_n-module M and submodule $N \subseteq M$, $d(M) = \max\{d(N), d(M/N)\}$.

Suppose $M_1, ..., M_k$ are all holonomic A_n-modules. Statement (b) follows by applying the above reasoning to the short exact sequence

$$0 \longrightarrow M_k \longrightarrow M_1 \oplus ... \oplus M_k \longrightarrow M_1 \oplus ... \oplus M_{k-1} \longrightarrow 0$$

and induction on k. □

Proposition 2.32. Holonomic modules are Artinian. Furthermore, their length is finite and bounded by their multiplicity.

Proof: Here we use the additivity of multiplicity from Proposition 2.27 (b). Let M be a holonomic left A_n-module and suppose we have a descending chain of proper submodules

$$M = N_0 \supsetneq N_1 \supsetneq N_2 \supsetneq \cdots \supsetneq N_k. \tag{$*$}$$

By Proposition 2.31, N_i and N_i/N_{i+1} are holonomic for each i. Together with the properness of the above inclusions, this implies $d(N_i) = d(N_i/N_{i+1}) = n$. We then have

$$m(M) = \sum_{i=0}^{k-1} m(N_i/N_{i+1}) \; + \; m(N_k).$$

Multiplicity is a nonnegative integer, and since the multiplicity of a nonzero module is by definition nonzero, $d(M) \geq k$ (allowing for the case that $N_k = 0$). However, $m(M)$ is itself a finite integer, so we cannot find a chain ($*$) of length greater than $m(M)$. In particular, any infinite chain must either stabilize, in which case $m(N_i/N_{i+1}) = 0$ for all $i \gg 0$, or terminate with $N_i = 0$ for all $i \gg 0$. □

2.3.4 Lemma on B-Functions

Let f be a polynomial in $K[x_1, ..., x_n]$ and let s be a new variable. We will consider the Weyl algebra $A_n(K(s))$ over the field of rational functions in s and the $A_n(K(s))$-module generated by the formal symbol f^s, upon which a rational function $p \in K(s)$ acts in the obvious way and the operator ∂_i acts by the formula

$$\partial_j\,(f^s) = \frac{s}{f} \cdot \frac{\partial f}{\partial x_i}. \tag{5}$$

Note that when we write f^{s+k} for some integer k, we mean $f^k \cdot f^s$. When s is an integer and f^s is treated not as a formal symbol but as a power, this action agrees with the existing action of ∂_j. The above formula means that $A_n(K(s))f^s$ is an $A_n(K(s))$-submodule of $K(s)[x_1, ..., x_n, f^{-1}]f^s$.

Lemma 2.33. Suppose M is a left A_n-module with a filtration Γ. If there exists a polynomial $q \in K[y]$ of degree n such that $\dim_K(\Gamma_i) \leq q(i)$ for sufficiently large i, then M is finitely generated and holonomic. In addition, if a is the leading coefficient of q, then $m(M) \leq n!a$.

Proof: Suppose first that $0 \neq N \subseteq M$ is a finitely generated submodule. We then have a good filtration Λ of N with respect to the Bernstein filtration on A_n by Proposition 2.21 as well as an induced filtration on N

given by $\Gamma_i \cap N$. By Proposition 2.24, there exists some positive integer k_0 such that $\Lambda_i \subseteq \Gamma_{i+k_0} \cap N$ for all $i \in \mathbb{N}$, and hence $\dim_K(\Lambda_i) \leq \dim_K(\Gamma_{i+k_0} \cap N) \leq q(i+k_0)$.

Let $\chi(t) = \chi(t, \Lambda, N)$ be the Hilbert polynomial for N with respect to Λ. For $i \gg 0$, we have

$$\chi(i) = \sum_{j=0}^{i} \dim_K(\Lambda_i/\Lambda_{i-1}) = \dim_K(\Gamma_i) \leq q(i+k_0).$$

This means $\deg(\chi) \leq \deg(q) = n$, and therefore N is holonomic by Bernstein's inequality. Since a polynomial converges to its largest term in the limit $t \longrightarrow \infty$, this also implies that $m(M) \leq n!a$, where a denotes the leading coefficient of q.

Consider now an ascending chain of finitely generated modules

$$N_0 \subsetneq N_1 \subsetneq N_2 \subsetneq \dots \subsetneq N_k$$

where $N_i \subseteq M$. Each of these is holonomic by what we have just proven. Repeating the argument from the proof of Proposition 2.32, we have that

$$m(N_k) = \sum_{i=1}^{k} m(N_i/N_{i-1}) \;+\; m(N_0) \geq k.$$

However, we also have that $m(N_k) \leq n!a$ by what we have already shown. This means that $n!a$ is an upper bound on the length on an ascending chain in M, and therefore M itself is finitely generated. Repeating the above argument for M, we get that $d(M) = n$ and $m(M) \leq n!a$. $\square$

The following corollary is crucial to the proof of Theorem 2.35.

Corollary 2.34. Fix a polynomial $f \in K[x_1, ..., x_n]$. The left $A_n(K(s))$-module $M = K(s)[x_1, ..., x_n, f^{-1}]f^s$ defined above is holonomic.

Proof: Let $m = \deg(f)$ in $K[x_1, ..., x_n]$. Define

$$\Gamma_k = \left\{ qf^{-k} \cdot f^s \;\middle|\; \deg(q) \leq (m+1)k \right\}.$$

We write $qf^{-k} \cdot f^s$ rather than qf^{s-k} to emphasize that qf^{-k} is an element in $K[x_1, ..., x_n f^{-1}]$ acting on f^s. Using the conventions of Definition 2.17, we show in detail that Γ is a filtration of M with $\dim_K(\Gamma_k) \leq \binom{n+k(m+1)}{k(m+1)}$. The holonomy of M then follows immediately from the previous lemma. Note that $\mathcal{B} = \{B_i\}_{i\in\mathbb{N}}$ is the Bernstein filtration on $A_n(K(s))$, as per usual.

(1) Clearly, if $qf^{-k} \cdot f^s \in \Gamma_k$, then

$$\deg(q \cdot f) = \deg(q) + \deg(f) \leq (m+1)k + m \leq mk + k + m + 1 = (m+1)(k+1).$$

Hence $qf^{-k} \cdot f^s = (qf)f^{-(k+1)} \cdot f^s \in \Gamma_{k+1}$, and therefore $\{\Gamma_i\}_{i\in\mathbb{N}}$ is an upward nested sequence of K-vector spaces.

(2) Fix $1 \leq i \leq n$. The left action of $\hat{x}_i \in A_n(K(s))$ on $qf^{-k} \cdot f^s \in \Gamma_k$ increases the degree of q by 1, so

$\hat{x}_i\left(qf^{-k}\cdot f^s\right)\in\Gamma_{k+1}$. The left action of ∂_{x_i} on $qf^{-k}\cdot f^s$ is given by

$$\partial_{x_i}(qf^{-k}\cdot f^s) = \partial_{x_i}(q)p^{-k}\cdot p^s \;-\; kp^{-(k+1)}\partial_{x_i}(f)q\cdot f^s + qf^s\frac{s}{f}f^s\partial_{x_i}(f)$$
$$= \left(\partial_{x_i}(q)f \;+\; (s-k)q\partial_{x_i}(f)\right)f^{-(k+1)}\cdot p^s.$$

Both terms inside the parentheses have degree at most $\deg(q)+m-1$, which is less than $(m+1)(k+1)$ because $\deg(q)\le(m+1)k$, so $\partial_{x_i}(qf^{-k}\cdot f^s)\in\Gamma_{k+1}$.

The set $\{\hat{x}_1,...,\hat{x}_n,\partial_{x_1},...\partial_{x_n}\}$ forms a basis for B_1, hence $B_1\cdot\Gamma_k\subseteq\Gamma_{k+1}$. Furthermore, $B_i\Gamma_k\subseteq\Gamma_{k+i}$ since $B_i=B^i$.

(3) Choose an arbitrary element $p\in K(s)[x_1,...,x_n,f^{-1}]$ so that $p\cdot f^s$ represents an arbitrary element of M. Set $k\le\deg(p)$ and $q=pf^k$. Then

$$p\cdot f^s = qf^{-k}\cdot f^s \;\text{ and }\; \deg(q)=\deg(f)+km\le k+km=(m+1)k,$$

so $f\cdot p^s\in\Gamma_k$. Every element of M is in Γ_k for some k, hence $\bigcup_{i=0}^{\infty}\Gamma_k=M$.

(4) The set of elements of the form $uf^{-k}\cdot f^s$ where u is a monomial of $K[x_1,...,x_n]$ with degree at most $(m+1)k$ generates Γ_k as a K-vector space, so each Γ_k is finite dimensional.

As discussed in the proof of Bernstein's theorem, there are $\binom{n+k(m+1)}{k(m+1)}$ many monomials in $K[x_1,...,x_n]$ of degree at most $(m+1)k$, so $\dim_K(\Gamma_k)\le\binom{n+k(m+1)}{k(m+1)}$ by (4) above. This binomial coefficient is a degree n polynomial in k, hence M is holonomic by Lemma 2.33. □

We can now prove the Lemma on b-functions. Like many other named lemmas in mathematics, it is listed not as a lemma but as a theorem.

Theorem 2.35. *Fix $f\in K[x_1,...,x_n]$. There exists a polynomial $B(s)\in K[s]$ and a differential operator $D(s)\in A_n(K)[s]$ such that*

$$B(s)f^s = D(s)f^{s+1}.$$

The set of all such $B(s)$ form an ideal in $K[s]$, the monic generator of which is called the Bernstein polynomial *of f and is denoted by $b_f(s)$.*

Proof: The case in which $f=0$ is trivial, so assume $f\neq 0$. Since $A_n(K(s))f^s$ is a submodules of $K(s)[x_1,...,x_nf^{-1}]f^s$, it too is holonomic and consequently of finite length. The descending sequence

$$A_n(K(s))\cdot f^s\supseteq A_n(K(s))\cdot f^{s+1}\supseteq A_n(K(s))\cdot f^{s+2}\supseteq ...$$

must therefore terminate. In particular, there must exist some positive integer k such that

$$A_n(K(s))f^k\cdot f^s = A_n(K(s))f^{k+1}\cdot f^s.$$

This implies that

$$f^{s+k}=D(s)f^{s+k+1}$$

for some $D(s) \in A_n(K(s))$. As s is simply a dummy variable, we can send $s \mapsto s-k$ to get $f^s = D(s-k)f^{s+1}$. Note that $D(s-k)$ is simply a polynomial in $\hat{x}_1, ..., \hat{x}_n, \partial_{x_1}, ..., \partial_{x_n}$ with coefficients in $K(s)$, so we may multiply by an appropriate $B(s) \in K[s]$ to clear denominators and get that $B(s)D \in A_n(K)[s]$. Setting $D'(s) = B(s)D(s-k)$ yields

$$B(s)f^s = D'(s)f^{s+1}$$

as desired. □

Example 2.36. Let $f = x_1^2 + ... + x_n^2$. Notice that

$$\partial_{x_i}^2 f^{s+1} = 4x_i^2(s+1)sf^{s-1} + 2(s+1)f^s.$$

Letting $D = \partial_{x_1}^2 + ... + \partial_{x_n}^2$, we get that

$$\begin{aligned} D(f^{s+1}) &= \sum_{i=0}^{n} \left(4x_i^2(s+1)sf^{s-1} \; + \; 2(s+1)f^s\right) \\ &= 4(s+1)s(x_1^2 + ... + x_n^2)f^{s-1} + 2n(s+1)f^s \\ &= 2(s+1)(2s+n)f^s, \end{aligned}$$

hence $b_f(s) = 2(s+1)(2s+n)f^s$.

A bountiful source of examples comes from Cayley's identity. This was proven by Cayley prior to the development of the theory of modules over the Weyl algebra.

Example 2.37. Let $f = \det(x_{ij})$ be the determinant of a $n \times n$ matrix (x_{ij}) in n^2 in indeterminates. Then

$$s(s+1)(s+2)...(s+n-1)f^s \;=\; \det(\partial/\partial x_{ij})f^{s+1}.$$

See [Ful14] for a proof of this fact. Let's check the 2×2 case. Let

$$A = \begin{pmatrix} x & y \\ z & w \end{pmatrix}, \qquad B = \begin{pmatrix} \partial_x & \partial_y \\ \partial_z & \partial_w \end{pmatrix}$$

and set $f = \det(A) = xw - yz$ and $D = \det(B) = \frac{\partial^2}{\partial x \partial w} - \frac{\partial^2}{\partial y \partial z}$. Then

$$\frac{\partial^2}{\partial x \partial w}(f^{s+1}) = (s+1)sf^{s-1}xw \qquad \text{and} \qquad \frac{\partial^2}{\partial y \partial z}(f^{s+1}) = (s+1)sf^{s-1}zy,$$

so

$$Df^{s+1} = s(s+1)f^{s-1}(xw - zy) = s(s+1)f^s.$$

2.4 Analogs for Algebraic *D*-Modules

We now wish to extend the results of this section to the setting of algebraic D-modules. This section is exclusively for reference; we prove almost nothing, and instead direct the reader to various sources. Throughout this section, X is a smooth variety over K and $\mathcal{M}$ is either a left or right D_X-module.

Just as in the affine case, we study D-modules through filtrations. The sheaf D_X is given locally on affines $\operatorname{Spec} R = U \subseteq X$ by operators $f \in D_R$, and likewise, for $k \in \mathbb{N}$ we may consider the coherent sheaf of $\mathcal{O}_X$-modules D_X^k given locally by order k operators $f \in D_R^k$. The collection $D_X^\bullet$ of coherent $\mathcal{O}_X$-modules is then a filtration of D_X. A filtration of $\mathcal{M}$ is then an increasing family of coherent submodules $\mathcal{F}_\bullet = \{\mathcal{F}_i\}_{i\in\mathbb{N}}$ each of which satisfies

$$D_X^i \cdot \mathcal{F}_j \subseteq \mathcal{F}_{i+j}.$$

This collection is also required to be exhausting, i.e.

$$\bigcup_{i\in\mathbb{Z}} \mathcal{F}_i(U) = \mathcal{M}(U) \quad \text{for each open } U \subseteq X.$$

We say that a filtration $\mathcal{F}_\bullet$ of $\mathcal{M}$ is a *good filtration* if the associated graded $\operatorname{gr}^{D_X^\bullet} D_X$-module

$$\operatorname{gr}^{\mathcal{F}_\bullet} \mathcal{M} \;=\; \bigoplus_{i\in\mathbb{Z}} \mathcal{F}_i/\mathcal{F}_{i-1}$$

is coherent.

Every coherent D_X-module $\mathcal{M}$ has a good filtration locally by Proposition 2.21. Somewhat more surprising is the fact that good filtrations are guaranteed to exist globally as well.

Lemma 2.38. Let $\mathcal{M}$ be a coherent D_X-module. Then there exists a good filtration $\mathcal{F}_\bullet$ of $\mathcal{M}$ by coherent $\mathcal{O}_X$-modules.

Proof: Let $\mathcal{M}$ be a left D_X-module, noting that the appropriate modifications to the following argument yield the same result when $\mathcal{M}$ is instead a right D_X-module. We will prove that there is a coherent $\mathcal{O}_X$-submodule $\mathcal{F} \subseteq \mathcal{M}$ which generates $\mathcal{M}$ over the action of D_X. The product $D_X^i\mathcal{F}$ will then be a coherent $\mathcal{O}_X$-module for each $i \in \mathbb{N}$ since both D_X^i and $\mathcal{F}$ are coherent $\mathcal{O}_X$-modules themselves. Defining

$$\mathcal{F}_i = D_X^i \cdot \mathcal{F} \subseteq \mathcal{M}$$

then gives a filtration $\mathcal{F}_\bullet$ of $\mathcal{M}$ by coherent $\mathcal{O}_X$-submodules.

Every variety is of finite type over its base field and is therefore quasi-compact. We can then find a finite cover $U_1, ..., U_n$ of X by open affine subsets, each of which is nonempty. Let S_i be a finite generating set for $\Gamma(U_i, \mathcal{M})$ over $\Gamma(U_i, D_X)$, which must exist by the coherence assumption on $\mathcal{M}$. The sheaf of $\mathcal{O}_{U_i}$-modules $\mathcal{F}_{U_i}$ defined to be the sheaf associated to the $\Gamma(U_i, \mathcal{O}_X)$-span of S_i is then a coherent by definition and generates $\mathcal{M}|_{U_i}$ over $D_X|_{U_i}$.

The trick, then, is to globalize these $\mathcal{O}_{U_i}$-submodules of $\mathcal{M}|_{U_i}$. Hartshorne [Har77, Exercise 2.5.15] gives us a method to do exactly that. Suppose X is a Noetherian scheme and $\mathcal{G}$ is a quasi-coherent $\mathcal{O}_X$-module. The exercise states that if $U \subseteq X$ is an open set and $\mathcal{F}_U$ is a coherent subsheaf of $\mathcal{G}|_U$, then there is a coherent subsheaf $\mathcal{F} \subseteq \mathcal{G}$ such that $\mathcal{F}|_U = \mathcal{F}_U$. As every variety is a Noetherian scheme, we get can find a coherent $\mathcal{O}_X$-submodule $\mathcal{F}_i \subseteq \mathcal{M}$ such that $\mathcal{F}_i|_{U_i} = \mathcal{F}_{U_i}$ for each $1 \leq i \leq n$. The universal property of sheafification then gives us a map

$$\mathcal{F}_1 \oplus ... \oplus \mathcal{F}_n \longrightarrow \mathcal{M},$$

whose image is a coherent $\mathcal{O}_X$-module which generates $\mathcal{M}$ as a D_X-module by construction. Hence we are done. $\square$

If $\mathcal{M}$ is a coherent D_X-module then we can find a good filtration $\mathcal{F}_\bullet$ of $\mathcal{M}$ by coherent $\mathcal{O}_X$-modules. The associated graded module $\operatorname{gr}^{\mathcal{F}_\bullet} \mathcal{M}$ is then coherent as a $\operatorname{gr}^{D_X^\bullet} D_X$-module. However, we no longer have access to the Bernstein filtration and its finite dimensionality, and therefore cannot apply the theory of Hilbert functions. We need another way to define dimension and introduce the following object.

Definition 2.39. The *characteristic variety* $\operatorname{Ch}(\mathcal{M})$ of $\mathcal{M}$ is the closed algebraic subset of T^*X given by $\widetilde{\operatorname{gr}^{\mathcal{F}_\bullet} \mathcal{M}}$, the sheaf associated to $\operatorname{gr}^{\mathcal{F}_\bullet} \mathcal{M}$, with reduced scheme structure.

The characteristic variety of a D-module, among other things, gives us another way to talk about dimension. Among other things, the dimension of the characteristic variety gives us an analog of Bernstein's inequality for modules over the Weyl algebra.

Proposition 2.40. Let $\mathcal{M}$ be a coherent left or right D_X-module. Then any irreducible component Λ of $\operatorname{Ch}(\mathcal{M})$ satisfies the inequality $\dim \Lambda \geq \dim X$. In particular, $\dim \operatorname{Ch}(\mathcal{M}) \geq \dim X$ if $\mathcal{M} \neq 0$.

Proof: This is a consequence of the following fact: $\operatorname{Ch}(\mathcal{M})$ is *involutive* with respect to the symplectic structure of T^*X. Sato-Kawai-Kashiwara originally proved this [SKK73] and Gabber later came up with algebraic proof [Gab81]. $\square$

Holonomy is therefore still a valid concept for D-modules over smooth varieties, and in fact the lemma on b-functions still holds. See [Kas77] for instance, or Mihnea Popa's online notes [Pop21].

3 Inverse Images, Direct Images and Kashiwara's Theorem

Given a morphism of smooth varieties $\varphi : X \longrightarrow Y$, we may push a sheaf on X forward via the direct image functor $\varphi_* : \mathrm{Sh}(X) \longrightarrow \mathrm{Sh}(Y)$ and pull a sheaf on Y back via the inverse image functor $\varphi^{-1} : \mathrm{Sh}(Y) \longrightarrow \mathrm{Sh}(X)$. These two operations are unfortunately not compatible with the action of differential operators; in general, the naïve pushforward and pullback of a D_X-module $\mathcal{M}$ along φ is not a D_Y-module.

Example 3.1. To see what goes wrong, let's examine the direct image of a left D_X-module $\mathcal{M}$ along the map $\varphi : \mathbb{A}^n_K \longrightarrow \mathbb{A}^n_k \times_K \mathbb{A}_K = \mathbb{A}^{n+1}_K$ defined by $x_{n+1} = 0$. This corresponds to the ring map $K[x_1, ..., x_{n+1}] \longrightarrow K[y_1, ..., y_n]$, $x_i \mapsto y_i$ for $1 \leq i \leq n$ and $x_{n+1} \mapsto 0$. There is a natural action of $\hat{x}_1, ..., \hat{x}_n, \hat{x}_{n+1}$ on $\varphi_*\mathcal{M}$, and letting ∂_{x_i} act via ∂_{y_i} for $1 \leq i \leq n$ causes no problems. However, since the action of x_{n+1} is trivial, for any section u of $\varphi_*\mathcal{M}$ we will always have

$$[\partial_{x_{n+1}}, \hat{x}_{n+1}]u = \partial_{x_{n+1}}(\hat{x}_{n+1}(u)) - \hat{x}_{n+1}(\partial_{x_{n+1}}(u)) = 0,$$

so there is no action of ∂_{n+1} will satisfy the relation $[\partial_{x_{n+1}}, \hat{x}_{n+1}] = 1$.

We therefore need to "fix" the direct and indirect image functors to work categories of D-modules. Unfortunately, a complete account of these topics requires the use of derived categories and Grothendieck's six functor formalism, topics beyond the scope of this essay. The problem is homological: the full statement of Kashiwara's theorem establishes an equivalence of categories via the direct image functor on the derived category of D-modules, but the candidates for this functor are not necessarily exact on the category of D-modules themselves.

Nonetheless, we can provide a meaningful discussion if we limit our discussion of direct images to closed embeddings $\iota : X \hookrightarrow Y$ and right D_X-modules. Our candidate definition for the direct images of D-module involves the left exact functor φ_* and the right exact functor $\otimes$, and is therefore neither left nor right exact itself. However, the direct image is exact when φ is a closed embedding. The inverse image will turn out to be only right exact.

As should be expected by now, many algebraic constructions are brushed under the rug. If the reader becomes stuck, the now familiar resources [Gin98] should prove quite useful, or [Har77] of course. The author additionally referenced [HTT08] and [Sch19] frequently while developing this section. Throughout, K is a field of characteristic zero and both X and Y are smooth algebraic varieties over K.

3.1 Inverse Images

Suppose $\varphi : X \longrightarrow Y$ is a morphism of smooth algebraic varieties over K and M is a left D_Y-module. We wish to build a left D_X-module from M in a meaningful way. The inverse image of M

$$\varphi^* M = \mathcal{O}_X \otimes_{\varphi^{-1}\mathcal{O}_Y} \varphi^{-1} M$$

is a left $\mathcal{O}_X$-module, and we can endow it with a left D_X-module structure as follows.

Fix a point $p \in Y$, an affine neighborhood U of p, a local coordinate system $\{y_i, \partial_{y_i}\}_{1 \leq i \leq n}$ of p on U, and set $V = \varphi^{-1}(U)$. It suffices to define the $\mathcal{O}_X(V)$ and $\Theta_X(V)$ action on elements of the form $r \otimes u \in \mathcal{O}_X(V) \otimes_{\varphi^{-1}\mathcal{O}_Y(V)} \varphi^{-1}M(V)$, as such elements generate $\varphi^{-1}M(V)$ and $\mathcal{O}_X$ and Θ_X generate D_X. We define the action of $a \in \mathcal{O}_X(V)$ on $r \otimes u$ by $a \cdot (r \otimes u) = ar \otimes u$ and the action of a vector field $\theta \in \Theta_X(V)$

on $r \otimes u$ by

$$\theta(r \otimes u) = \theta(r) \otimes u \;+\; r \sum_{i=1}^{n} \theta(y_i \circ \varphi) \otimes \partial_{y_i}(u). \tag{$*$}$$

To check that this does indeed produce a D_X-action on $\varphi^* M$, we need to verify that it satisfies the relations

$$\begin{aligned} [\partial_{x_i}, \hat{x}_j] &= \delta_{ij} \\ [\hat{x}_i, \hat{x}_j] &= [\partial_{x_i}, \partial_{x_j}] = 0 \end{aligned}$$

in an affine neighborhood $U' \subseteq X$ of $\varphi^{-1}(p)$ with a local coordinate system $\{x_i, \partial_{x_i}\}_{1 \le i \le m}$. We check the first relation on and claim the others follows similarly. For $r \otimes u \in \varphi^* M$, we have

$$\begin{aligned} \partial_{x_i} \hat{x}_j (r \otimes u) &= \partial_{x_i}(x_j r \otimes u) \\ &= \partial_{x_i}(x_j r) \otimes u \;+\; x_j r \sum_{k=1}^{n} \partial_{x_i}(y_k \circ \varphi) \otimes \partial_{y_k}(u) \\ &= r\delta_{ij} \otimes u \;+\; x_j \partial_{x_i}(r) \otimes u \;+\; x_j r \sum_{k=1}^{n} \partial_{x_i}(y_k \circ \varphi) \otimes \partial_{y_k}(u) \\ &= \delta_{ij}(r \otimes u) \;+\; x_j \left(\partial_{x_i}(r) \otimes u \;+\; r \sum_{k=1}^{n} \partial_{x_i}(y_k \circ \varphi) \otimes \partial_{y_k}(u) \right) \\ &= \delta_{ij}(r \otimes u) \;+\; \hat{x}_j \partial_{x_i}(r \otimes u), \end{aligned}$$

hence $[\partial_{x_i}, \hat{x}_j](r \otimes u) = \delta_{ij}(r \otimes u)$. It holds on arbitrary elements of $\varphi^* M$ by the linearity of the commutator.

This discussion is summarized by the following definition.

Definition 3.2. Let $\varphi : X \longrightarrow Y$ be a morphism of smooth algebraic varieties and let M be a D_Y-module. Then the inverse image $\varphi^* M$ of M endowed with the action defined in ($*$) is D_X-module, the *inverse image* of M.

Remark 3.3. While φ^{-1} is exact, the functor $\mathcal{O}_X \otimes_{\varphi^{-1}\mathcal{O}_Y} -$ is only right exact in general. This means φ^* is also only right exact. To preserve homological data, it is typical to work in the derived setting and replace φ^* with its left derived functor $\mathcal{O}_X \otimes^{L}_{\varphi^{-1}\mathcal{O}_Y} -$. The definition provided will be suitable for our needs, however.

As a sanity check, let's ensure the inverse image works as expected when φ is the identity map.

Example 3.4. Let $\varphi : X \longrightarrow X$ be the identity morphism on a smooth variety X and M a D_X-module. Note that the presheaf $U \mapsto \mathcal{O}_X(U) \otimes_{\mathcal{O}_X(U)} M(U)$ is a sheaf. We have $\varphi^{-1}(\mathcal{F})(U) = \mathcal{F}(U)$ for any sheaf $\mathcal{F}$ on X since φ is the identity, hence for any open set $V \subseteq X$,

$$\varphi^* M(V) = \mathcal{O}_X(V) \otimes_{\varphi^{-1}\mathcal{O}_X(V)} \varphi^{-1} M(V) \cong \mathcal{O}_X(V) \otimes_{\mathcal{O}_X(V)} M(V) \cong M(V).$$

Fix a point $p \in X$, an affine open neighborhood $U \subseteq X$ of p, and a local coordinate system $\{x_i, \partial_{x_i}\}_{1 \le i \le n}$ at p on U. Let $\theta \in \Theta_X(U)$ be a vector field on U and let $\theta = \sum_{i=1}^{n} a_i \partial_{x_i}$ be θ expressed in local coordinates (here,

$a_i \in \mathcal{O}_X(U)$). For any $u \in M$, we have that

$$\begin{aligned}\theta(1 \otimes u) &= \theta(1) \otimes u \;+\; \sum_{i=1}^{n} \theta(x_i \circ \varphi) \otimes \partial_{x_i}(u) \\ &= \sum_{i=1}^{n} \theta(x_i) \otimes \partial_{x_i}(u) \\ &= \sum_{i=1}^{n} a_i \otimes \partial_{x_i}(u) \\ &= 1 \otimes \left(\sum_{i=1}^{n} a_i \partial_{x_i}(u) \right) = 1 \otimes \theta(u),\end{aligned}$$

so $\varphi^* M \cong M$ via the isomorphism $1 \otimes u \mapsto u$.

However, inverse images can behave badly even for relatively simple morphisms $\varphi : X \longrightarrow Y$. For instance, the inverse image of a coherent module need not itself be coherent, as seen in the following example.

Example 3.5 (Loss of Coherence)**.** Suppose $X = Y = \mathbb{A}^1_K$, so that $D_X = D_Y = \tilde{A}_1$, the sheaf associated to the first Weyl algebra. Though X and Y are two copies of the same variety, we distinguish the coordinate systems of X and Y by $\{x, \partial_x\}$ and $\{y, \partial_x\}$ respectively, noting that these are globally valid.

Consider the morphism $\varphi : X \longrightarrow Y$ defined $\varphi(x) = x^2$ and note that the induced map on global sections $\varphi^\sharp : K[y] \longrightarrow K[x]$ sends a polynomial $f(y)$ to $f(x^2)$. Finally, let $M = A_1$, so that $\tilde{M}$ is Weyl algebra considered as a module over itself.

Hartshorne tells us that $\varphi^*(M) \cong (K[x] \otimes_{K[y]} M)^\sim$ [Har77, Proposition 2.5.2], so the global sections of $\varphi^*(M)$ are generated by elements of the form $f \otimes u$ for $f \in K[x]$ and $u \in M$. Though $\tilde{M}$ is coherent as a D_Y-module, we will see that $\varphi^* \tilde{M}$ is not a coherent D_X module.

It suffices to check that $\Gamma(X, \varphi^*(\tilde{M})) = K[x] \otimes_{K[y]} M$ is not finitely generated as a $\Gamma(X, D_X) = A_1$-module. Suppose we have some finite set of elements $B \subseteq K[x] \otimes_{K[y]} M$. The span of an element $f \otimes u \;+\; f' \otimes u'$ is contained in the span of $\{f \otimes u, f' \otimes u'\}$, so we may assume that B is comprised entirely of elements $f \otimes u$ for $f \in K[x]$ and $u \in M$. Furthermore, by writing u in its canonical form (see Lemma 1.17) we may assume that $u = \hat{y}^a \partial_y^b$ for some $a \in \mathbb{N}$ and $b \in \mathbb{N}$.

Suppose b is the largest natural number such that $f \otimes \hat{y}^a \partial_y^b$ is an element of B for some $a \in \mathbb{N}$ and $f \in K[x]$. From the $K[y]$-action on $K[x]$, we get that $f \otimes \hat{y}^a \partial_y^b = x^{2a} f \otimes \partial_y^b$. Noting that $x \circ \varphi = x^2$, we have

$$\begin{aligned}\partial_x(f \otimes \hat{y}^a \partial_y^b) &= \partial_x(x^{2a} f \otimes \partial_y^b) \\ &= \partial(x^{2a} f) \otimes \partial_y^b \;+\; x^{2a} f \partial_x(x^2) \otimes \partial_y(\partial_y^b) \\ &= \left(2a x^{2a-1} f(x) + x^{2a} f'\right) \otimes \partial_y^b \;+\; 2x^{2a+1} f \otimes \partial_y^{b+1}.\end{aligned}$$

Thus, the action of ∂_x will increase the degree of both the first and second component of $x^{2a} f \otimes \partial_y^b$ by 1. This means the A_1-span of $K[x] \otimes_{K[y]} M$ avoids elements such as $1 \otimes \partial_y^{b+1}$, as 1 has degree 0 and ∂_y^{b+1} has degree larger than b, the largest power of ∂_y appearing in the set B. Therefore, the span of any finite subset of $K[x] \otimes_{K[y]} M$ will be a proper subset, so $\varphi^*(\tilde{M})$ is not a coherent D_X-module.

There are a class of morphisms which *do* preserve finite generation of modules over the Weyl algebra under pullbacks, namely projections.

Example 3.6 (Projections)**.** Let $X = \mathbb{A}_K^m$ and $Y = \mathbb{A}_K^n$, so that $X \times_K Y \cong \mathbb{A}_K^{m+n}$. We denote by $K[X,Y]$ the polynomial ring $K[x_1, ..., x_m y_1, ..., y_n]$. Suppose that M is a left A_n-module and $\pi : X \times_K Y \longrightarrow Y$ is the projection map defined by $x_1 = ... = x_m = 0$. By [Har77, Proposition 2.5.2],

$$\pi^*(M) \cong K[X,Y] \otimes_{K[Y]} M$$

and this isomorphism can be checked to be compatible with the action of differential operators. As a K-vector space, $\pi^*(M)$ is generated by elements of the form $pq \otimes u$, where p is a monomial in $x_1, ..., x_m$, q is a monomial in $y_1, ..., y_n$. However, $pq \otimes u = p \otimes qu$, so we have an isomorphism

$$\pi^*(M) \cong K[X] \otimes_K M$$

of K-vector spaces. Because $K[X]$ is a left A_m-module and M is a left A_n-module, $K[Y] \otimes_K M$ has the natural structure of a A_{m+n}-module. That is, an operator $pq\partial^\alpha\partial^\beta$ in A_{m+n} with $p \in K[X]$, $q \in K[Y]$, $\alpha \in \mathbb{N}^m$ and $\beta \in \mathbb{N}^m$ acts on $f \otimes u \in K[X] \otimes$ by

$$pq\partial^\alpha\partial^\beta(f \otimes u) = p\partial^\alpha(f) \otimes q\partial^\beta(u).$$

We show that this coincides with the one defined by definition 3.2.

The action of polynomials in $K[X,Y]$ is clear. It then suffices to check the action of ∂_{x_j} and ∂_{y_k} on a generator $pq \otimes u$, where $1 \leq j \leq m$, $1 \leq k \leq n$, p is a monomial in $K[X]$ and q is a monomial in $K[Y]$. Noting that $y_i \circ \pi = y_i$ for each $1 \leq i \leq n$, we have

$$\begin{aligned}
\partial_{x_j}(pq \otimes u) &= \partial_{x_j}(pq) \otimes u \ + \ pq \sum_{i=1}^n \partial_{x_j}(y_i) \otimes \partial_{y_i} u \\
&= q\partial_{x_j}(p) \otimes u \\
&= \partial_{x_j}(p) \otimes qu
\end{aligned}$$

and

$$\begin{aligned}
\partial_{y_k}(pq \otimes u) &= \partial_{y_k}(pq) \otimes u \ + \ pq \sum_{i=1}^n \partial_{y_k}(y_i) \otimes \partial_{y_i}(u) \\
&= p\partial_{y_k}(q) \otimes u \ + \ pq \otimes \partial_{y_k}(u) \\
&= p \otimes \partial_{y_k}(qu),
\end{aligned}$$

where the last equality follows from Leibniz's rule. Hence, $\pi^*(M) \cong K[X] \otimes_K M$ as a A_{m+n}-module under the identification $K[X,Y] \otimes_{K[Y]} M \cong K[X] \otimes_K M$. In particular, since $K[X]$ is finitely generated as a A_m-module, if M is finitely generated as a A_n-module then $\pi^*(M)$ is finitely generated as a A_{m+n}-module.

3.2 Direct Images

The inverse image of D_Y is the module $\varphi^* D_Y = \mathcal{O}_X \otimes_{\varphi^{-1}\mathcal{O}_Y} \varphi^{-1} D_Y$. This inverse image is special, for in addition to the left D_X-action discussed in the previous section, it comes equipped with a right $\varphi^{-1} D_Y$ action. These actions are compatible, and therefore $\varphi^* D_Y$ is a $(D_X, \varphi^{-1} D_Y)$-bimodule. We give is a special name.

Definition 3.7. Suppose $\varphi : X \longrightarrow Y$ is a morphism of smooth varieties. We define the *transfer module* $D_{X \to Y}$ to be the $(D_X, \varphi^{-1} D_Y)$-bimodule $\varphi^* D_Y = \mathcal{O}_X \otimes_{\varphi^{-1}\mathcal{O}_Y} \varphi^{-1} D_Y$.

Let $\iota : X \longrightarrow Y$ be a closed embedding of smooth varieties and recall from the last section that the transfer module $D_{X\to Y} = \iota^* D_Y$ is a $(D_X, \iota^{-1}D_Y)$-bimodule. Let's examine $D_{X\to Y}$ in local coordinates.

Lemma 3.8. The transfer module $D_{X\to Y}$ is a locally free D_X-module of infinite rank and contains a copy of D_X

Proof: Fix an affine open $V \subseteq Y$, then $U = \iota^{-1}(V)$ is an open affine in X. Every smooth variety is locally a complete intersection; hence, we can choose local coordinates $y_1, ..., y_n$ for Y on V such that X is defined by the equations $y_{r+1} = ... = y_n = 0$. Let $\partial_{y_1}, ..., \partial_{y_n}$ be the vector fields corresponding to the y_i for Y on V and let $\partial_{x_1}, ..., \partial_{x_r}$ be the vector fields corresponding to $x_1 = y_1, ..., x_r = y_r$ for X on U (see [HTT08, Theorem A.5.3], for instance).

The transfer module has a global section $1 \otimes \hat{1}$, where 1 is the identity section of $\mathcal{O}_X$ and $\hat{1}$ is the section corresponding to the identity operator in $\iota^{-1}D_Y$. The actions of D_X on $1 \otimes \hat{1}$ gives us an embedding $D_X \longrightarrow D_{X\to Y}$, which can be easily seen in local coordinates. The action of ∂_{x_i} on $1 \otimes 1$ as defined in $(*)$ from the previous section is

$$\partial_{x_i}(1 \otimes \hat{1}) = \partial_{x_i}(1) \otimes \hat{1} \; + \; \sum_{j=1}^{n} \partial_{x_i}(y_j \circ \iota) \otimes \partial_{y_j} \circ \hat{1} = 1 \otimes \partial_{y_i}.$$

More generally, an operator $P = \sum_\alpha f_\alpha \partial^\alpha \in D_X(U)$ with $\alpha \in \mathbb{N}^r$ acts on $1 \otimes \hat{1}$ by the formula

$$P(1 \otimes \hat{1}) = \sum_\alpha f_\alpha \otimes \partial_{y_1}^{\alpha_1} ... \partial_{y_r}^{\alpha_r}.$$

Hence $P(1 \otimes 1) = 0$ only if $P = 0$ in $D_X(U)$, so the map $D_X \longrightarrow D_{X\to Y}$ given by the action of D_X on $1 \otimes \hat{1}$ is injective.

We now show $D_{X\to Y}$ is a locally free left D_X-module. Define the following subalgebra of D_Y:

$$D_Y^X = \bigoplus_{\alpha\in\mathbb{N}^r} \mathcal{O}_Y \cdot \partial_{y_1}^{\alpha_1} \cdot ... \cdot \partial_{y_r}^{\alpha_r}.$$

This is identical to D_Y itself except that we only allow vector fields which act nontrivially on the image of X in Y. Once we add the rest of the vector fields back in we recover D_Y, i.e. the map $D_Y^X \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}] \longrightarrow D_Y$ given by multiplication is an isomorphism. Furthermore, by the discussion above, we see that the map $D_X \longrightarrow D_{X\to Y}$ identifies D_X with the subalgebra $\mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}D_Y^X$ of $D_{X\to Y}$. But then

$$\begin{aligned} D_{X\to Y} &= \mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}D_Y \\ &\cong \mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}\left(D_Y^X \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}]\right) \\ &\cong \left(\mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}D_Y^X\right) \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}] \\ &\cong D_X \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}], \end{aligned}$$

where we use the identification $D_Y^X \cong \mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}D_Y^X$ in the final isomorphism. This implies $D_{X\to Y}$ is locally a free D_X-module. □

Remark 3.9. Notice that the above proof implies that inverse images over closed embeddings do not necessarily preserve coherence. Though D_Y is certainly locally finitely generated as a module over itself, its inverse image $D_{X\to Y}$ is locally a free module of infinite rank.

The functor $- \otimes_{D_X} D_{X\to Y}$ is exact on left and right D_X-modules since $D_{X\to Y}$ is locally free. The direct image functor on $\mathcal{O}_X$-modules ι_* is also exact since ι is a closed embedding, so $\iota_*(- \otimes_{D_X} D_{X\to Y})$ is exact. Thus the following definition fits the desired criterion for a direct image functor on right D_X-modules.

Definition 3.10. Let $\mathcal{M}$ be a right D_X-module. The *direct image* or *pushforward* of $\mathcal{M}$ along ι is

$$\iota_+ M = \iota_*(M \otimes_{D_X} D_{X\to Y}).$$

This is a right D_Y-module under the morphism $D_Y \longrightarrow \iota_*\iota^{-1}D_Y$.

Remark 3.11. When ι is replaced by an arbitrary map $\varphi : X \longrightarrow Y$, φ_* is only left exact and $- \otimes_{D_X} D_{X\to Y}$ is only right exact. We can fix this by replacing φ_* with its right derived functor $R\varphi_*$ and $- \otimes_{D_X} D_{X\to Y}$ with the left derived tensor product $- \otimes^L_{D_X} D_{X\to Y}$ [HTT08].

Remark 3.12. The inclusion of $D_X \hookrightarrow D_{X\to Y}$ given by Lemma 3.8 induces a similar inclusion of $\iota_*\mathcal{M}$ into $\iota_+\mathcal{M}$. On an affine open $U \subseteq Y$, we have

$$\iota_+\mathcal{M} \cong \iota_* M \otimes_K K[\partial_{y_r+1}, ..., \partial_{y_n}]$$

in the local coordinates defined in the proof of the lemma. The pushforward as we've defined it therefore solves the issue of D_Y's action on $\iota_*\mathcal{M}$ by simply attaching a copy of $\iota_*\mathcal{M}$ to each monomial in $\partial_{y_{r+1}}, ..., \partial_{y_n}$. We can also see that the copy of $\iota_*\mathcal{M}$ given by $\iota_*\mathcal{M} \otimes \hat{1}$ is exactly the submodule of $\iota_+\mathcal{M}$ annihilated by the ideal sheaf $\mathcal{I}_X \subseteq \mathcal{O}_Y$ defined by the closed embedding. This is clear locally, since $\mathcal{I}_Y(U) = (y_{r+1}, ..., y_n) \subseteq \mathcal{O}_Y(U)$.

Example 3.13. Let's compute the direct image of D_X along $\iota : X \hookrightarrow Y$. We get

$$\iota_+ D_X = \iota_* \left(D_X \otimes_{D_X} D_{X\to Y}\right) = \iota_* D_{X\to Y}.$$

The sheaf $i_* D_{X\to Y}$ is $i_*(\mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}D_Y)$. But we have a natural map $D_Y \longrightarrow i_*(\mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}D_Y)$ defined on sections by $P \mapsto 1 \otimes P$, and this is surjective. Furthermore, a section P is sent to 0 in $i_* D_{X\to Y}$ if and only if it can be written locally on an affine U as a sum with coefficients in $\mathcal{I}_X(U) \subseteq \mathcal{O}_Y(U)$, so the kernel of the map $D_Y \longrightarrow \iota_+ D_X$ is $\mathcal{I}_X D_Y$, the sheaf of ideals given by the kernel of $\iota^\sharp : \mathcal{O}_Y \longrightarrow \iota_*\mathcal{O}_X$. Hence

$$\iota_+ D_X \cong i_*(\mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}D_Y) \cong D_Y/\mathcal{I}_X D_Y.$$

Example 3.14. Now let us discuss the case of the embedding $\iota : \mathbb{A}^r_K \hookrightarrow \mathbb{A}^n_K$ from the beginning of this section, noting that because we're working over affine varieties we can perform all computations on global sections. Denote the corresponding quotient map by $\varphi : K[y_1, ..., y_n] \longrightarrow K[x_1, ..., x_r]$ and let A_r and A_n be the rth and nth Weyl algebras respectively. We compute the direct image of the right A_r-module $M = A_r/(P_1, ..., P_m)A_r$. We can realize M as the cokernel of the map

$$A_r^{\oplus m} \xrightarrow{(P_1,...,P_m)} A_r,$$

and because ι_+ is exact, it preserves this presentation. From the last example, $\iota_+ A_r \cong A_n/\ker\varphi A_n = A_n/(y_{r+1}, ..., y_n)A_n$, so $\iota_+ M$ is the cokernel of the map

$$(A_n/(y_{r+1}, ..., y_n)A_n)^{\oplus m} \longrightarrow A_n/(y_{r+1}, ..., y_n)A_n$$

induced by ι_+. If $\alpha_P : A_r \longrightarrow A_r$ is the morphism given by multiplication on the right by $P \in A_r$, then the induced endomorphism of $A_n/(y_{r+1}, ..., y_n)A_n$ is still multiplication on the right by P, albeit with P expressed in the coordinates $y_1, ..., y_r, \partial_{y_1}, ..., \partial_{y_r}$. Hence $\iota_+ M \cong A_n/(P_1, ..., P_r, y_{r+1}, ..., y_n)$.

3.2.1 Direct Images for Affines

In the case that $X = \operatorname{Spec} A$, $Y = \operatorname{Spec} B$ and $\varphi : X \longrightarrow Y$ is any regular map, the pushforward $\varphi_*(\tilde{M})$ of a quasi-coherent $\mathcal{O}_X$-module $\tilde{M}$ is simply the restriction of scalars along the ring homomorphism $\varphi^\sharp : B \longrightarrow A$ corresponding to φ. This leaves the underlying Abelian group of M unchanged and is hence exact. This means that when X and Y are affine, replacing ι with φ Definition 3.10 gives us a reasonable definition for the pushforward along φ.

We also have a way of obtaining left modules from right modules in the affine case, as described in Example 2.7. Using [Har77, Proposition 2.5.2] gives us that $D_{X\to Y} = \varphi^*(D_B) \cong (A \otimes_B D_B)$, so

$$D_{Y\leftarrow X} := {}^t(D_{X\to Y}) \cong {}^t(D_B) \otimes_B {}^t(A)$$

is a (D_B, D_A)-bimodule by right action of D_A on A. The rightmost isomorphism follows from the fact that ${}^t(M_1 \otimes_R M_2) \cong {}^t(M_2) \otimes_R {}^t(M_1)$, which can be quickly seen from definitions or found in [Cou95, Chapter 16].

Definition 3.15. Let $\varphi : X \longrightarrow Y$ be a morphism of smooth affine varieties and $\mathcal{M}$ be a left D_X-module. The *direct image* of $\mathcal{M}$ by φ is the left D_Y-module

$$\varphi_+ M = D_{Y\leftarrow X} \otimes_{D_X} M.$$

Let's use this to compute the pushforward of left and right $D_{X\times_K Y}$-modules along the projection onto Y when $X = \mathbb{A}^m_K$ and $Y = \mathbb{A}^n_K$.

Example 3.16. Let $\pi : X \times_K Y \longrightarrow Y$ be the projection onto Y and M be a right A_{m+n}-module. By Example 3.6 we know

$$D_{X\to Y} = \pi^* A_n = K[X] \otimes_K A_n,$$

and since $K[X] \cong A_m/\sum_{i=1}^m A_m\partial_{x_i}$, we have

$$D_{X\to Y} = A_m \Big/ \sum_{i=1}^m A_m \partial_{x_i} \otimes_K A_n \cong A_{m+n} \Big/ \sum_{i=1}^m A_{m+n}\partial_{x_i}.$$

Since $\sum_{i=1}^m A_{m+n}\partial_{x_i}$ is a (A_{m+n}, A_n)-bimodule, the isomorphism

$$\pi_+ M = M \otimes_{A_m+n} \left(A_{m+n} \Big/ \sum_{i=1}^m A_{m+n}\partial_{x_i} \right) \cong M \Big/ \sum_{i=1}^m M\partial_{x_i}$$

of A_n-modules then follows from basic properties of tensor products. If N is instead a left A_{m+n}-module, then we simply need to compute $D_{Y\leftarrow X}$. The standard anti-automorphism of A_{m+n} takes $\partial_{x_i} \mapsto -\partial_{x_i}$, so

$$D_{Y\leftarrow X} = {}^t\!\left(A_{m+n}/\sum_{i=1}^m A_{m+n}\partial_{x_i} \right) \cong A_{m+n} \Big/ {}^t\!\left(\sum_{i=1}^m A_{m+n}\partial_{x_i} \right) \cong A_{m+n} \Big/ \sum_{i=1}^m \partial_{x_i} A_{m+n}.$$

The fact ${}^t(A_n/J) = A_n/{}^t(J)$ used above follows immediately from considering the map $A_n \longrightarrow {}^t(A_n) \longrightarrow {}^t(A_n/J)$. It then follows that

$$\pi^* N \cong A_{m+n} \Big/ \sum_{i=1}^{m} \partial_{x_i} A_{m+n} \ \otimes_{A_{m+n}} \ N \cong N \Big/ \sum_{i=1}^{m} \partial_{x_i} N.$$

Note that the quotient $A_{m+n}/\sum_{i=1}^m A_{m+n}\partial_{x_i}$ appearing throughout this example is not a A_{m+n}-module since the left ideal $\sum_{i=1}^m A_{m+n}\partial_{x_i}$ is not two-sided. Nevertheless, it is a A_n module under the embedding $A_n \hookrightarrow A_{m+n}$ because the generators of A_n all commute with the ∂_{x_i}'s.

3.3 Kashiwara's Equivalence

In the name of further developing our functorial language, we establish some categorical notation standard to the literature.

Definition 3.17 (Notation)**.** We establish the following notation.

- $\mathsf{Mod}_{\mathrm{qc}}(D_X)$ is the category of left D_X-modules and $\mathsf{Mod}_{\mathrm{c}}(D_X)$ is the category of coherent left D_X-modules.
- If $X \subseteq Y$ is a closed embedding, then ${\mathsf{Mod}_{\mathrm{qc}}}^X(D_Y)$ is the category of left D_Y-modules with support in X, likewise ${\mathsf{Mod}_{\mathrm{c}}}^X(D_Y)$ is the category of coherent left D_Y modules with support in X.

The category of right D_X-modules can be identified with $\mathsf{Mod}_{\mathrm{qc}}(D_X^{\mathrm{op}})$, so we obtain the right module versions of the above definitions by replacing D_X with D_X^{op}.

Let $\iota : X \longrightarrow Y$ once again be a closed embedding. It is convenient to identify X with its image in Y, and indeed we can easily do so under the identifications $x_i \mapsto y_i$ in local coordinates as given in the beginning of the proof of Lemma 3.8. We can now state a version of Kashiwara's equivalence theorem.

Theorem 3.18. *Let $\iota : Y \hookrightarrow X$ be a closed embedding. The functor ι_+ induces the following equivalences of categories*

$$\iota_+ : \mathsf{Mod}_{qc}(D_X^{\mathrm{op}}) \longrightarrow {\mathsf{Mod}_{qc}}^X(D_Y^{\mathrm{op}}),$$
$$\iota_+ : \mathsf{Mod}_{c}(D_X^{\mathrm{op}}) \longrightarrow {\mathsf{Mod}_{c}}^X(D_Y^{\mathrm{op}})$$

between right D_X-modules and right D_X-modules with

Proof: Let $\mathcal{M}$ be a right D_Y-module. This proof naturally splits into three parts.

The first step is to construct a functor $\iota^\natural$ which will serve as the inverse to ι_+. We saw in Remark 3.12 that $\iota_*\mathcal{M}$ embeds into $\iota_+\mathcal{M}$ and is exactly the subsheaf annihilated by $\mathcal{I}_X \subseteq \mathcal{O}_Y$. The functor $\iota^\natural$ should therefore take a D_Y-module $\mathcal{N}$ to its subsheaf of sections annihilated by $\mathcal{I}_X$, somehow interpreted as a D_X-module.

There is a technical, yet efficient way of doing this. Given a right D_Y-module $\mathcal{N}$, define

$$\iota^\natural\mathcal{N} = \mathcal{H}om_{\iota^{-1}D_Y}(D_{X\to Y}, \iota^{-1}\mathcal{N}).$$

This is a right $\iota^{-1}D_Y$-module by definition and has a right D_X-module structure induced by the left D_X

action on $D_{X\to Y}$. We may rewrite this as

$$\begin{aligned}\iota^\natural \mathcal{N} &= \mathcal{H}om_{\iota^{-1}D_Y}(\mathcal{O}_X \otimes_{\iota^{-1}\mathcal{O}_Y} \iota^{-1}D_Y, \iota^{-1}\mathcal{N}) \\ &\cong \mathcal{H}om_{\iota^{-1}\mathcal{O}_Y}\left(\mathcal{O}_X, \mathcal{H}om_{\iota^{-1}D_Y}(\iota^{-1}D_Y, \iota^{-1}\mathcal{N})\right) \\ &\cong \mathcal{H}om_{\iota^{-1}\mathcal{O}_Y}(\mathcal{O}_X, \iota^{-1}\mathcal{N})\end{aligned}$$

using Hom-tensor adjunction. By definition of $\mathcal{I}_X$ we have a short exact sequence

$$0 \longrightarrow \mathcal{I}_X \longrightarrow \mathcal{O}_Y \longrightarrow \iota_*\mathcal{O}_X \longrightarrow 0.$$

Applying the exact functor ι^{-1} gives us

$$0 \longrightarrow \iota^{-1}\mathcal{I}_X \longrightarrow \iota^{-1}\mathcal{O}_Y \longrightarrow \mathcal{O}_X \longrightarrow 0$$

and applying $\mathcal{H}om_{\iota^{-1}\mathcal{O}_Y}(-, \iota^{-1}\mathcal{N})$ gives us

$$0 \longrightarrow \iota^\natural\mathcal{N} \longrightarrow \iota^{-1}\mathcal{N} \longrightarrow \mathcal{H}om_{\iota^{-1}\mathcal{O}_Y}(\iota^{-1}\mathcal{I}_X, \iota^{-1}\mathcal{N}).$$

The rightmost map can be factored as the isomorphism $\iota^{-1}\mathcal{N} \longrightarrow \mathcal{H}om_{\iota^{-1}\mathcal{O}_Y}(\iota^{-1}\mathcal{O}_Y, \iota^{-1}\mathcal{N})$ and the restriction $\mathcal{H}om_{\iota^{-1}\mathcal{O}_Y}(\iota^{-1}\mathcal{O}_Y, \iota^{-1}\mathcal{N}) \longrightarrow \mathcal{H}om_{\iota^{-1}\mathcal{O}_Y}(\iota^{-1}\mathcal{I}_X, \iota^{-1}\mathcal{N})$. The kernel of this map is locally the sections of $\iota^{-1}\mathcal{N}$ which are annihilated by $\iota^{-1}\mathcal{I}_X$, so $\iota^\natural\mathcal{N}$ is exactly the subsheaf of $\iota^{-1}\mathcal{N}$ annihilated by $\iota^{-1}\mathcal{I}_X$ by the exactness above. This isomorphism can be seen to be compatible with the natural D_X actions on the respective sheaves, but we omit this detail.

In the second part of the proof, we must prove that $\iota^\natural$ and ι_+ are indeed inverses. That is, for a D_X-module $\mathcal{M}$ and a D_Y-module $\mathcal{N}$ supported on X, we must show that the natural morphisms $\iota^\natural\iota_+\mathcal{M} \longrightarrow \mathcal{M}$ and $\mathcal{N} \longrightarrow \iota_+\iota^\natural\mathcal{N}$ are isomorphisms. It suffices to check this locally, so we assume without loss of generality that $Y = \operatorname{Spec} B$ is affine with local coordinate system $\{y_i, \partial_{y_i}\}$ and that X is defined by the ideal $(y_{r+1}, ..., y_n) \subseteq B$. Setting $A = B/I$, we then have $X = \operatorname{Spec} B/I$. In this local setting, the pushforward of a right D_A-module M is then exactly $M \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}]$ as in Remark 3.12. This submodule of ι_+M annihilated by I is then $M \otimes 1 \cong M$, proving the first isomorphism.

The second isomorphism requires more work. Fix a right D_B-module N such that $\operatorname{Supp}(N) \subseteq V(I) = X$ so that every element of N is annihilated by a sufficiently high power of I. Set $N_0 = \{u \in N \mid uI = 0\}$k, and note that we can consider N_0 to be a right D_A-module under the identifications $x_i \mapsto y_i$ and $\partial_{x_i} \mapsto \partial_{y_i}$. Our goal then is to show $N \cong \iota_+N_0 = N_0 \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}]$. The key to this will be the action of the operator $T_j = \hat{y}_j\partial_{y_j}$ on N for $r+1 \leq j \leq n$. The point is this: T_j acts trivially on N_0 by definition and

$$T_j \cdot \partial_{y_{r+1}}^{e_{r+1}} \cdot ... \cdot \partial_{y_n}^{e_n} = \partial_{y_{r+1}}^{e_{r+1}} \cdot ... \cdot \partial_{y_n}^{e_n}(T_j - e_j),$$

which means

$$u \otimes \partial_{y_{r+1}}^{e_{r+1}} \cdot ... \cdot \partial_{y_n}^{e_n} \cdot (T_j - e_j) = 0.$$

Elements of the form $u \otimes \partial_{y_{r+1}}^{e_{r+1}} \cdot ... \cdot \partial_{y_n}^{e_n}$ are therefore eigenvalues of T_j with corresponding eigenvalue e_j, at least when T_j is considered to be an operator on $N_0 \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}]$. Such elements also form a basis for $N_0 \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}]$ over K, so every eigenvector is of this form.

We use this intuition to construct an isomorphism $N \cong N_0 \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}]$ by considering a decomposition of N into eigenspaces. One can show by expanding that

$$T_j(T_j - 1)...(T_j - e) = \hat{y}_{y_j}^{e+1} \partial_{y_j}^{e+1} \quad \text{for } \ e \geq \mathbb{N}.$$

Since each $u \in N$ is annihilated by a sufficiently large power of x_j,

$$u \cdot T_j(T_j - 1)...(T_j - e) = \hat{y}_{y_j}^{e+1} \partial_{y_j}^{e+1} = \quad \text{for } \ e \gg 0.$$

In particular, this implies u can be written as a sum of eigenvectors of T_j with eigenvalues in $\mathbb{N}$. The operators $T_{r+1}, ..., T_n$ all commute, so we obtain a decomposition

$$N \cong \bigoplus_{\eta \in \mathbb{N}^{n-r}} N_\eta$$

of N into simultaneous eigenstates, where T_j acts on N_η via multiplication on the right by η_j. If $s \in N_\eta$, then $s\partial_{y_j} \in N_{\eta + e_j}$ since T_j acts via multiplication by $\eta_j + 1$:

$$s\partial_{y_j} T_j = s(\partial_{y_j} \hat{y}_j)\partial_{y_j} = s(\hat{y}_j \partial_{y_j} + 1)\partial_{y_j} = s\partial_{y_j}(\eta_j + 1).$$

Similarly, $s\hat{y}_j \in N_{e-1}$ since T_j acts via multiplication by $\eta_j - 1$:

$$s\hat{y}_j T_j = s\hat{y}_j(\hat{y}_j \partial_{y_j}) = s\hat{y}_j(\partial_{y_j} \hat{y}_j - 1) = s\hat{y}_j(\eta_j - 1).$$

Since N_η is trivial whenever η has a component less than 0, $N_{(0,...,)}$ must be killed by $\hat{y}_{r+1}, ..., \hat{y}_n$, hence $N_{(0,...,0)} = N_0$. Furthermore, since each T_j commutes with $\hat{y}_1, ..., \hat{y}_r, \partial_{y_1}, \partial_{y_r}$, each N_η is a D_A-module and the maps

$$N_0 \longrightarrow N_\eta, \ s \mapsto s\partial_{y_{r+1}}^{\eta_{r+1}} ... \partial_{y_n}^{\eta_n}$$

are all isomorphisms of D_A-modules. We conclude that the map

$$N_0 \otimes_K K[\partial_{y_{r+1}}, ..., \partial_n] \longrightarrow N, \quad \sum u \otimes \partial_{y_{r+1}}^{\eta_{r+1}} ... \partial_{y_n}^{\eta_n} = \sum u\partial_{y_{r+1}}^{\eta_{r+1}} ... \partial_{y_n}^{\eta_n}$$

is an isomorphism, as desired.

For the final part of this proof, we must check that ι_+ and $\iota^\natural$ preserve coherence. It also suffices to check this locally, so we may assume $Y = \operatorname{Spec} B$ and $X = \operatorname{Spec} A$ as above. Assume N is a finitely generated right D_B-module with support in X. By what we have just shown, $N \cong N_0 \otimes_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}]$, and hence is generated as a D_B-module by finitely many elements $s_1, ..., s_k \in N_0$. But then $\iota^\natural N = \iota^\natural N_0$ is generated as a D_A-module by $s_1, ..., s_k$.

Likewise, if M is a finitely generated right D_A-module, then $\iota_+ M \cong_K K[\partial_{y_{r+1}}, ..., \partial_{y_n}]$ is finitely generated as a D_B-module, and we are done. □

3.4 Preservation of Holonomy

Our final, brief discussion is one that will tie the topics of dimension, holonomy, filtrations, functoriality, and Fourier transforms together. We have seen that neither inverse images nor direct images over arbitrary maps of varieties $\varphi : X \longrightarrow Y$ preserve coherence, but somewhat surprisingly, they *do* preserve holonomy.

It turns out that the proof of this statement for smooth varieties can be reduced to the case that $X = \mathbb{A}_K^{n+1}$, $Y = \mathbb{A}_K^n$, and φ is the projection map defined by $x_{n+1} = 0$. To see why the justification for this in full detail, see [HTT08, Chapter 3]. Briefly, this is accomplished by decomposing an arbitrary regular map $\varphi : X \longrightarrow Y$ as the composition of a closed embedding $\iota : X \longrightarrow X \times_K Y$ and a projection $X \times_K Y$. It then suffices to assume φ is either a closed embedding or a projection. The closed embedding case is handled fairly quickly by [HTT08, Lemma 3.2.5]. When φ is a projection, the problem is further reduced to the special case mentioned above. The proof of this final case fits nicely into the theory we developed over the course of this essay, and we present it now as our final result.

Let us first state some lemmas.

Lemma 3.19. Suppose M is a finitely generated A_n-module and N is a finitely generated A_m-module. Then

(a) $d(M \otimes_K N) = d(M) \ + \ d(N)$,

(b) $m(M \otimes_K N) = m(M) \cdot m(N)$.

Proof: [Cou95, Theorem 13.4.1]. □

Lemma 3.20. If $\iota : X \longrightarrow Y$ is a closed embedding of smooth varieties and $\mathcal{M}$ is a holonomic D_Y-module, then $\iota^*\mathcal{M}$ is a holonomic D_X-module.

Proof: [HTT08, Theorem 3.2.3 and Lemma 3.2.5]. □

Lemma 3.21. Let M be a finitely generated left A_n-module and let $\widehat{M}$ be the Fourier transform of M. Then $d(M) = d(\widehat{M})$ and $m(M) = m(\widehat{M})$.

Proof: Let $\mathcal{B} = \{B_k\}_{k \in \mathbb{N}}$ be the Bernstein filtration and F be the automorphism of A_n defined $F(\hat{x}_i) = -\partial_i$ and $F(\partial_i) = \hat{x}_i$ for $1 \leq i \leq n$. This preserves the degree of elements in A_n, and hence $F(B_k) = B_k$ for all $k \in \mathbb{N}$.

As in Proposition 2.21, let Γ be a good filtration for M such that Γ_0 is a finite dimensional K-vector space whose basis is a set of generators for M. Then $\widehat{M}$ is also generated by Γ_0 over A_n. Defining $\Omega_k = F(B_k) \cdot \Gamma_0$ thus defines a good filtration Ω on $\widehat{M}$, and since $F(B_k) \cdot \Gamma_0 = B_k\Gamma_0 = \Gamma_k$, we have $\Omega = \Gamma$. This implies that $\operatorname{gr}^\Gamma M$ and $\operatorname{gr}^\Omega \widehat{M}$ have the same Hilbert polynomial, and hence M and $\widehat{M}$ have the same dimension and multiplicity. □

We can now prove that pushforwards and pullbacks of holonomic modules over the projection $\mathbb{A}_K^{n+1} \longrightarrow \mathbb{A}_K^n$ are themselves holonomic.

Theorem 3.22. *Let $X = \mathbb{A}_K^1$, $Y = \mathbb{A}_K^n$, and $\pi : X \times_K Y \longrightarrow Y$ be the projection defined by $x = 0$. If $\mathcal{M}$ is a holonomic left $D_{X \times_K Y}$ module and $\mathcal{N}$ is a holonomic left D_Y-module, then $\pi_+\mathcal{M}$ is a holonomic D_Y-module and $\pi^*\mathcal{N}$ is a holonomic $D_{X \times_K Y}$-module.*

Proof: We can once again work with global sections since all varieties are affine. Set $M = \Gamma(X \times_K Y \mathcal{M})$ and $N = \Gamma(Y, \mathcal{N})$ so that M and N are holonomic modules over A_{n+1} and A_n respectively.

Let us first show that π^*N is holonomic as a A_{n+1}-module. From Example 3.6, we know $\pi^*N \cong K[x] \otimes_K N$, and by Lemma 3.19 above, $d(\pi^*N) = d(K[x]) + d(N) = 1 + n$. Hence π^*N is a holonomic A_{n+1}-module.

Now consider the pushforward $\pi_+ M$. By Example 3.16 this is isomorphic to $M/\partial_x M$ with the action of A_n given by the embedding $A_n \hookrightarrow A_{n+1}$. The same arguments used in Example 2.3 show this is isomorphic to the Fourier transform of M/xM, and therefore M/xM and $\pi_+ M$ have the same dimension by Lemma 3.21. We will write M/xM as the pullback of M along a closed embedding, which is necessarily holonomic by Lemma 3.20.

Consider the closed embedding $\iota : Y \hookrightarrow X \times_K Y$ corresponding to the surjective ring homomorphism $K[x, y_1, ..., y_n] \longrightarrow K[y_1, ..., y_n]$. Since we are working entirely over affines, the pullback of M is simply

$$\iota^* M = K[y_1, ..., y_n] \otimes_{K[x,y_1,...,y_n]} M$$

by definition, which is isomorphic to M/xM as a $K[x, y_1, ..., y_n]$-module. For an element $q \otimes u \in \iota^* M$ and a derivation ∂_{y_i} of A_n, we have

$$\begin{aligned} \partial_{y_i}(q \otimes u) &= \partial_{y_i}(q) \otimes u \ + \ q \sum_{k=1}^{n} \frac{\partial y_k}{\partial y_i} \otimes \partial_{y_k}(u) \\ &= \partial_{y_i}(q) \otimes u \ + \ q \otimes \partial_{y_i}(u), \end{aligned}$$

which corresponds to the element $\partial_{y_i}(qu)$ in M/xM, so the isomorphism $\iota^* M = M/xM$ is compatible with the relations of A_n. Hence M/xM and consequently $\pi_+ M$ are holonomic. $\square$

References

[AM16] M. F. Atiyah and I. G. Macdonald. *Introduction to commutative algebra*. economy. Addison-Wesley Series in Mathematics. For the 1969 original see [MR0242802]. Westview Press, Boulder, CO, 2016, pp. ix+128. ISBN: 978-0-8133-5018-9; 0-201-00361-9; 0-201-40751-5.

[Ber] Joseph Bernstein. *Algebraic Theory of D-modules*. Online Lecture Notes. URL: https://ncatlab.org/nlab/files/BernsteinDModule.pdf.

[Ber72] I. N. Bernstein. "Analytic continuation of generalized functions with respect to a parameter". In: *Funkcional. Anal. i Prilozen.* 6.4 (1972), pp. 26–40. ISSN: 0374-1990.

[BH93] Winfried Bruns and Jürgen Herzog. *Cohen-Macaulay rings*. Vol. 39. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge, 1993, pp. xii+403. ISBN: 0-521-41068-1.

[Cou95] S. C. Coutinho. *A primer of algebraic D-modules*. Vol. 33. London Mathematical Society Student Texts. Cambridge University Press, Cambridge, 1995, pp. xii+207. ISBN: 0-521-55119-6; 0-521-55908-1. DOI: 10.1017/CBO9780511623653. URL: https://doi.org/10.1017/CBO9780511623653.

[Ful14] Markus Fulmek. "A combinatorial proof for Cayley's identity". In: *Electron. J. Combin.* 21.4 (2014), Paper 4.40, 17.

[Gab81] Ofer Gabber. "The integrability of the characteristic variety". In: *Amer. J. Math.* 103.3 (1981), pp. 445–468. ISSN: 0002-9327. DOI: 10.2307/2374101. URL: https://doi.org/10.2307/2374101.

[Gin98] Victor Ginzburg. *Lectures on D-modules*. Online Lecture Notes. 1998. URL: https://people.math.harvard.edu/~gaitsgde/grad_2009/Ginzburg.pdf.

[Har77] Robin Hartshorne. *Algebraic geometry*. Graduate Texts in Mathematics, No. 52. Springer-Verlag, New York-Heidelberg, 1977, pp. xvi+496. ISBN: 0-387-90244-9.

[HTT08] Ryoshi Hotta, Kiyoshi Takeuchi, and Toshiyuki Tanisaki. *D-modules, perverse sheaves, and representation theory*. Vol. 236. Progress in Mathematics. Translated from the 1995 Japanese edition by Takeuchi. Birkhäuser Boston, Inc., Boston, MA, 2008, pp. xii+407. ISBN: 978-0-8176-4363-8. DOI: 10.1007/978-0-8176-4523-6. URL: https://doi.org/10.1007/978-0-8176-4523-6.

[Jef20] Jack Jeffries. *D-modules*. Online Lecture Notes. 2020. URL: https://jack-jeffries.github.io/S2020/D-modules.pdf.

[Kas77] Masaki Kashiwara. "B-functions and holonomic systems. Rationality of roots of B-functions". In: *Invent. Math.* 38.1 (1976/77), pp. 33–53. ISSN: 0020-9910. DOI: 10.1007/BF01390168. URL: https://doi.org/10.1007/BF01390168.

[Kas97] Masaki Kashiwara. "Erratum: "Algebraic study of systems of partial differential equations" [Mém. Soc. Math. France (N.S.) No. 63 (1995); MR1384226 (97f:32012)]". In: *Bull. Soc. Math. France* 125.2 (1997), p. 313. ISSN: 0037-9484. URL: http://www.numdam.org/item?id=BSMF_1997__125_2_313_0.

[Moo04] Rahim Moosa. *Jet spaces in complex analytic geometry: an exposition*. 2004. DOI: 10.48550/ARXIV.MATH/0405563. URL: https://arxiv.org/abs/math/0405563.

[MR01] J. C. McConnell and J. C. Robson. *Noncommutative Noetherian rings*. Revised. Vol. 30. Graduate Studies in Mathematics. With the cooperation of L. W. Small. American Mathematical Society, Providence, RI, 2001, pp. xx+636. ISBN: 0-8218-2169-5. DOI: 10.1090/gsm/030. URL: https://doi.org/10.1090/gsm/030.

[Muh88] Jerry L. Muhasky. "The differential operator ring of an affine curve". In: *Trans. Amer. Math. Soc.* 307.2 (1988), pp. 705–723. ISSN: 0002-9947. DOI: 10.2307/2001194. URL: https://doi.org/10.2307/2001194.

[Pop21] Mihnea Popa. *The Bernstein Sato Polynomial*. Online Lecture Notes. 2021. URL: https://people.math.harvard.edu/~mpopa/notes/Bernstein-Sato-notes.pdf.

[Sch19] Christian Schnell. *Algebraic D-Modules*. Online Lecture Notes. 2019. URL: http://www.math.stonybrook.edu/~cschnell/mat615/.

[SKK73] Mikio Sato, Takahiro Kawai, and Masaki Kashiwara. "Microfunctions and pseudo-differential equations". In: *Hyperfunctions and pseudo-differential equations (Proc. Conf., Katata, 1971; dedicated to the memory of André Martineau)*. 1973, 265–529. Lecture Notes in Math., Vol. 287.