The Riemann-Hilbert correspondence

or: I only know how to work with curves

Sanath Devalapurkar

March 22, 2020

Outline

Hilbert's twenty-first problem

Consider the differential equation

$$z\frac{df}{dz} = \lambda f$$

with $\lambda \in \mathbf{C}$, on $\mathbf{C} \setminus \{0\}$. It is solved by $f = z^\lambda$, and looping around the origin (via the path $z \mapsto e^{2\pi i t}z$) sends f to $e^{2\pi i \lambda}f$. This defines a monodromy representation

$$\begin{aligned}\pi_1(\mathbf{C} \setminus \{0\}) = \mathbf{Z} &\to \mathrm{GL}_1(\mathbf{C}) \\ 1 &\mapsto e^{2\pi i \lambda}.\end{aligned}$$

One might ask for the converse: given a representation ρ of the fundamental group ("monodromy"), can one find a linear differential equation whose monodromy is ρ?

If one further allows for certain singularities in the differential equation, this is Hilbert's twenty-first problem.

Fancification

Since we're all fancy and modern, we know that differential equations are just connections on vector bundles. Indeed, a differential equation

$$\frac{d^n f}{dz^n} + p_1(z)\frac{d^{n-1} f}{dz^{n-1}} + \cdots + p_n(z) f = 0$$

defined on $\mathbf{A}^1$, say, is just the bundle $\mathcal{O}_{\mathbf{A}^1}^{\oplus n}$ equipped with the connection

$$\nabla : \mathcal{O}_{\mathbf{A}^1}^{\oplus n} \to \mathcal{O}_{\mathbf{A}^1}^{\oplus n} \otimes_{\mathcal{O}_{\mathbf{A}^1}} \Omega^1_{\mathbf{A}^1} = (\Omega^1_{\mathbf{A}^1})^{\oplus n}$$

sending $\vec{f}$ to $d\vec{f} + A \cdot \vec{f} dz$, where

$$A = \begin{pmatrix} 0 & -1 & 0 & \cdots & 0 \\ 0 & 0 & -1 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & -1 \\ p_n(z) & \cdots & \cdots & \cdots & p_1(z) \end{pmatrix}.$$

Fancification, part 2

We also know that local systems are just representations of π_1.
So, as a baby case, we'd like to understand the relationship between:

- local systems;
- vector bundles equipped with a connection.

Here's a theorem.

Riemann-Hilbert

Let X be a complex manifold. Then there is an equivalence of categories:

$$\{\text{Local systems on } X\} \xrightarrow{\sim} \begin{Bmatrix} \text{Vector bundles on } X + \\ \text{flat/integrable connection} \end{Bmatrix};$$

$$\mathcal{L} \mapsto (\mathcal{L} \otimes_{\mathbf{C}} \mathcal{O}_X, 1 \otimes d);$$

$$\ker(\nabla) \leftarrow\!\shortmid (\mathcal{F}, \nabla)$$

It's easy to see that if $\mathcal{L}$ is a local system, then $(\mathcal{L} \otimes_{\mathbf{C}} \mathcal{O}_X, 1 \otimes d)$ is a vector bundle with flat connection: $d^2 = 0$.

Why's it true?

The other direction is more subtle. If $(\mathcal{F}, \nabla)$ is a vector bundle with flat connection, you need:

- $\ker(\nabla)$ to be a local system;
- $\mathcal{F} \cong \ker(\nabla) \otimes_{\mathbf{C}} \mathcal{O}_X$.

The second is a consequence of flatness (move along the fibers).

The first is a jazzed up version of the local uniqueness of first-order ODEs with an initial condition:
Write $\nabla = d + A$ with A a matrix of 1-forms. Then $f = (f_1, \cdots, f_n) \in \ker(\nabla)$ iff

$$df_i + \sum_j A_{ij} f_j = 0.$$

By uniqueness of solutions, you can find a unique solution for every point in $\mathbf{C}^n$ (corresponding to an initial value of f), and so locally $\ker(\nabla) \cong \mathbf{C}^n$.

How do we generalize?

We'd like to generalize in two directions:

- go to the algebraic setting;
- allow more general D-modules.

Roadblock

Consider $X = \mathbf{A}^1_{\mathbf{C}} = \operatorname{Spec} \mathbf{C}[z]$. Look at two connections on $\mathcal{O}_X$:

$$\nabla(f) = df,\ \nabla'(f) = df + fdz.$$

Same analytically: $f \mapsto fe^z$. So the horizontal sections on X^{an} are the same. Look at the differential equation associated to ∇':

$$f' + f = 0.$$

Formally define $t = z^{-1}$; then, this becomes

$$-t^2 f'(t) + f(t) = 0.$$

This has an *irregular singularity* at $t = 0$, i.e., at $z = \infty$.

18.03 stuff(??)

Consider a differential equation

$$\frac{d^n f}{dz^n} + p_1(z)\frac{d^{n-1} f}{dz^{n-1}} + \cdots + p_n(z) f = 0,$$

with each p_i meromorphic.

Definition

This differential equation has a *regular singularity* at z_0 if:

- not all p_k are analytic at z_0;
- all $(z - z_0)^k p_k(z)$ are analytic at z_0, i.e., p_k has a pole of order at most k at z_0.

If the second condition is not satisfied, this differential equation is said to have an *irregular singularity* at z_0.

Our discussion implies that we should impose some regularity conditions on our D-modules. We'll return to this soon.

More restrictions: holonomicity

Regularity won't suffice: we need an analogue of integrability/flatness.
Flatness got us uniqueness of solutions to ODEs, which was crucial to the solutions/horizontal sections being a local system.
The appropriate generalization of flatness is called *holonomicity*, and that's what we'll now talk about.

The definition

Holonomic D-modules are "maximally overdetermined" systems of differential equations.

Consider $X = \mathbf{A}^n$, so $\mathcal{D}_X = \mathbf{C}[x_1, \cdots, x_n]\langle\partial_1, \cdots, \partial_n\rangle$ is the Weyl algebra. Suppose I is a left $\mathcal{D}_X$-ideal. For instance, if $n = 1$ and $I = \langle\partial_x\rangle$, then $\mathcal{D}_X/I = \mathcal{O}_X$.

Obtain the characteristic variety $\mathrm{Ch}(I)$ of I in $\mathbf{A}^{2n} = \mathrm{T}^*X$: look at zeros of the principal symbols $\sigma(P)$ for all $P \in I$.
A theorem of Bernstein's implies that $\dim \mathrm{Ch}(I) \geq n$.

Definition

The $\mathcal{D}_X$-module $\mathcal{D}_X/I$ is said to be *holonomic* if $\dim \mathrm{Ch}(I) = n$.

Examples

Suppose $n = 1$, and say $I = \mathcal{D}_X \cdot P$, where

$$P = \sum_{i=0}^{n} p_i(z)\partial_z^i.$$

Then $\sigma(P) = \sum_{i=0}^{n} p_i(z)t^i$, where (z, t) are the coordinates in $\mathrm{T}^*\mathbf{A}^1 = \mathbf{A}^2$. So if P is nontrivial, then $\mathcal{D}_X/\mathcal{D}_X \cdot P$ is holonomic.

But if $n > 1$, and I is generated by a single differential operator, will *never* be holonomic.
However, if I is generated by n distinct equations, then $\mathcal{D}_X/\mathcal{D}_X \cdot P$ will be holonomic if $\mathrm{Ch}(\mathcal{F})$ is a complete intersection.

The $\mathcal{D}_X$-module $\mathcal{O}_X$ is holonomic, but $\mathcal{D}_X$ is not.

The general definition

We would like to define holonomicity on a general smooth $\mathbf{C}$-variety X.

Recall that there is a filtration $\mathcal{D}_X^{\leq n}$ by the order of a differential operator, such that $\mathrm{gr}(\mathcal{D}_X) = \mathcal{O}_{\mathrm{T}^*X}$.

Good filtration

A filtration $F_i\mathcal{F}$ on a $\mathcal{D}_X$-module $\mathcal{F}$ is *good* if $\mathrm{gr}(\mathcal{F})$ is a coherent $\mathcal{O}_{\mathrm{T}^*X}$-module.

The characteristic variety $\mathrm{Ch}(\mathcal{F})$ is $\mathrm{Supp}(\mathrm{gr}(\mathcal{F})) \subseteq \mathrm{T}^*X$.

Turns out this is independent of the choice of good filtration.
A theorem of Bernstein's says that $\dim(\mathrm{Ch}(\mathcal{F})) \geq \dim(X)$.

Holonomicity

A $\mathcal{D}_X$-module is said to be *holonomic* if $\dim(\mathrm{Ch}(\mathcal{F})) = \dim(X)$.

We can extend to complexes $\mathcal{F} \in \mathrm{D}^b(\mathcal{D}_X)$ by asking that each cohomology sheaf be holonomic. Write $\mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_X)$ for the full subcategory on holonomic complexes.

More examples

As before, $\mathcal{O}_X$ is holonomic. So is $\mathcal{D}_X/\mathcal{D}_X \cdot (z\partial_z - \lambda) = \mathcal{O}_X\{z^\lambda\}$, as well as $\mathcal{D}_X/\mathcal{D}_X \cdot (\partial_z - 1) = \mathcal{O}_X\{e^z\}$.

Again, $\mathcal{D}_X$ is not holonomic.

A filtration $F_i\mathcal{F}$ being good turns out to be equivalent to asking that:

- Each $F_i\mathcal{F}$ is coherent;
- for $i \gg 0$, we have $F_1\mathcal{D}_X \cdot F_i\mathcal{F} = F_{i+1}\mathcal{F}$.

If $\mathcal{F}$ is coherent over $\mathcal{O}_X$ (so it's a vector bundle with a flat connection), then take $F_i\mathcal{F} = \mathcal{F}$. This is obviously a good filtration, and $\mathrm{gr}(\mathcal{F}) = \mathcal{F}$.
So $\mathrm{Ch}(\mathcal{F}) \subseteq X$, and Bernstein implies that $\mathrm{Ch}(\mathcal{F}) = X$. Therefore:

Observation

Vector bundles with flat connections are holonomic.

Holonomicity is not too bad

Holonomic D-modules are "generically vector bundles with flat connections":

Theorem

Let $\mathcal{F}$ be a holonomic $\mathcal{D}_X$-module. Then there exists an open dense $U \subseteq X$ such that $\mathcal{F}|_U$ is coherent over $\mathcal{O}_U$.

Idea: $\mathrm{gr}(\mathcal{F})|_X$ is coherent (because $\mathcal{F}$ is holonomic); so let $\mathrm{gr}(\mathcal{F})^0$ be the submodule supported away from $X \subseteq \mathrm{T}^*X$. If $\pi : \mathrm{T}^*X \to X$, let $U = X \setminus \pi(\mathsf{Supp}(\mathrm{gr}(\mathcal{F})^0))$.

In fact, it's really nice!

Theorem

Let $\mathcal{F} \in \mathrm{D}^b(\mathcal{D}_X)$. TFAE:

- $\mathcal{F}$ is holonomic;
- for each $i_x : \{x\} \hookrightarrow X$, the complex $i_x^\dagger \mathcal{F}$ is finite-dimensional;
- there is a finite sequence

$$\emptyset = X_{n+1} \subseteq X_n \subseteq \cdots \subseteq X_0 = X$$

of closed subsets with $X_r \setminus X_{r+1}$ smooth, and each $i_r^\dagger \mathcal{F}$ a coherent $\mathcal{O}_X$-module, where $i_r : X_r \setminus X_{r+1} \hookrightarrow X$.

Philosophical takeaway

This theorem says that holonomic $\mathcal{D}_X$-modules $\mathcal{F}$ are *precisely* those for which there exist a stratification over which $\mathcal{F}$ corresponds to a flat connection on each stratum.

Comparing to the Riemann-Hilbert correspondence discussed above, we would expect that a "solution sheaf" that might be associated to $\mathcal{F}$ should be such that there exists a stratification on which the solution sheaf is locally constant. This is precisely a constructible sheaf on X (we'll define this later).

But we aren't done yet, because we still need to impose some regularity conditions (as we saw in the beginning).

One more theorem

In the previous talk, we defined a bunch of functors on D-modules. Holonomicity behaves well with respect to all of them:

Theorem

The duality functor **D** induces an equivalence $\mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_X) \xrightarrow{\sim} \mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_X)^{\mathrm{op}}$. Moreover, the external tensor product induces functors

$$\boxtimes : \mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_X) \times \mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_Y) \to \mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_{X \times Y}).$$

If $f : X \to Y$ is a morphism of smooth algebraic varieties, then

$$\int_f, \int_{f!}, \ f^\dagger, f^\star : \mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_X) \to \mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_Y).$$

Motivation

We now turn to regularity. We gave a definition above, but we'll have to rephrase it.

Again consider a differential equation

$$\frac{d^n f}{dz^n} + p_1(z)\frac{d^{n-1} f}{dz^{n-1}} + \cdots + p_n(z) f = 0,$$

with each p_i meromorphic.
This is equivalent to a differential equation

$$b_n(z)\theta^n + b_{n-1}(z)\theta^{n-1} f + \cdots + b_0(z) f = 0,$$

with each b_i meromorphic, where $\theta = z\partial_z$.

Motivation, continued

In turn, we get a system of ODEs

$$\frac{d}{dz}\vec{f}(z) = \frac{\Gamma(z)}{z}\vec{f}(z), \tag{1}$$

where

$$\Gamma(z) = \begin{pmatrix} 0 & -1 & 0 & \cdots & 0 \\ 0 & 0 & -1 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & -1 \\ b_0/b_n & b_1/b_n & b_2/b_n & \cdots & b_{n-1}/b_n \end{pmatrix}.$$

The function f is a solution to the original ODE iff $(f, \theta f, \cdots, \theta^{n-1} f)$ is a solution to (1).

Theorem (Fuchs)

Our ODE is regular if and only if each b_i/b_n is holomorphic.

Regularity on the disk

This motivates a possible approach to defining regularity.

Let's stay in the analytic world: let D be the unit complex disk, and $\mathring{D}$ the punctured disk. So $\mathcal{O}_D[z^{-1}]$ consists of meromorphic functions on D which are holomorphic on $\mathring{D}$.

A meromorphic connection ∇ on a vector bundle $\mathcal{F}$ over D is a morphism $\nabla : \mathcal{F} \to \Omega^1_D[z^{-1}] \otimes_{\mathcal{O}_D} \mathcal{F}$ satisfying the Leibniz rule $\nabla(fs) = df \otimes s + f \cdot \nabla(s)$.

Definition

Say that ∇ is regular if there is a choice of local coordinates $e_1, \cdots, e_n$ on $\mathcal{F}$ such that

$$\nabla e_i = \sum_j \frac{b_{ij}(z)}{z} e_j,$$

with each $b_{ij}(z)$ holomorphic.

In other words, if we locally write $\nabla = d + A$ with A a matrix of sections of $\Omega^1_D[z^{-1}]$, then the entries of A have poles of order at most 1 at $z = 0$ (up to a possible gauge transformation).

Now, notice that if $(\mathcal{F}, \nabla)$ is a regular meromorphic connection, then for any $f \in \mathcal{F}$, there is a $\mathcal{O}_D$-finitely generated submodule

$$\mathcal{L} = \mathcal{O}_D f + \mathcal{O}_D \theta f + \cdots + \mathcal{O}_D \theta^{n-1} f \subseteq \mathcal{F}$$

such that $f \in \mathcal{L}$ and $\theta \mathcal{L} \subseteq \mathcal{L}$.

Conversely, if for any $f \in \mathcal{F}$, there is a $\mathcal{O}_D$-finitely generated submodule $\mathcal{L} \subseteq \mathcal{F}$ such that $f \in \mathcal{L}$ and $\theta \mathcal{L} \subseteq \mathcal{L}$, then define

$$\mathcal{L}_i = \mathcal{O}_D f + \mathcal{O}_D \theta f + \cdots + \mathcal{O}_D \theta^{i-1} f.$$

Finite generation implies that $\mathcal{L}_{m+1} = \mathcal{L}_m$ for some $m \gg 0$. This forces

$$\theta^m f = -\sum_{i=0}^{m-1} b_i \theta^i f.$$

So:

Observation

$(\mathcal{F}, \nabla)$ is a regular meromorphic connection iff $\mathcal{F}$ is a union of θ-stable $\mathcal{O}_D$-finitely generated submodules.

Considerations in dimension one

Now we're cooking!
Let $X = \mathbf{A}^1$ and $U = \mathbf{A}^1 \setminus \{0\}$, so $i : U \hookrightarrow X$.
Note:

$$\mathcal{D}_X = \mathbf{C}\langle z, \partial_z \rangle,\ \mathcal{D}_U = \mathbf{C}\langle z, z^{-1}, \partial_z \rangle.$$

Let $\theta = z\partial_z$, and let $\mathcal{D}_X^0 = \mathbf{C}\langle z, \theta \rangle$.

Definition

A $\mathcal{O}_U$-coherent $\mathcal{D}_U$-module $\mathcal{F}$ is said to be *regular at zero* if $\int_i \mathcal{F}$ is a union of $\mathcal{O}_X$-finitely generated $\mathcal{D}_X^0$-submodules.

Let's unpack this.

Examples

- Let
$$\mathcal{F} = \mathcal{D}_U/\mathcal{D}_U \cdot \partial_z = \mathcal{O}_U = \mathbf{C}[z^{\pm 1}].$$
Then $\int_i \mathcal{F} = \mathbf{C}[z^{\pm 1}]$, and $\theta(z^n) = nz^n$. Because $\mathbf{C}[z^{\pm 1}] = \bigcup_n z^{-n}\mathbf{C}[z]$, and $z^{-n}\mathbf{C}[z]$ is a $\mathcal{D}_X^0$-submodule with one generator, we conclude that $\mathcal{F}$ is regular.
- More generally, let
$$\mathcal{F} = \mathcal{D}_U/\mathcal{D}_U \cdot (z\partial_z - \lambda) = \mathcal{O}_U\{z^\lambda\}.$$
Then θ acts on z^λ by λz^λ. The same argument shows that it's regular.
- Let
$$\mathcal{F} = \mathcal{O}_U\{\log(z)\}.$$
Then θ acts on $\log(z)$ by sending it to 1. Since $\int_i \mathcal{F} = \bigcup_n z^{-n}\mathbf{C}[z]\log(z)$, again find that $\mathcal{F}$ is regular at zero.

A non-example

For a non-example, let

$$\mathcal{F} = \mathcal{D}_U/\mathcal{D}_U \cdot (z^{n+1}\partial_z + n) = \mathcal{O}_U\{\exp(z^{-n})\}$$

for $n > 0$. Then θ acts on $\exp(z^{-n})$ by $-nz^{-n}\exp(z^{-n})$. It follows that

$$\mathcal{D}_X^0\{\exp(z^{-n})\} = \mathbf{C}\langle z, \theta\rangle\{\exp(z^{-n})\} = \int_i \mathcal{F}.$$

The issue is that when we differentiated, the degree dropped by two.

Note that *analytically*, $\mathcal{F}$ is the same as $\mathcal{O}_U = \mathcal{D}_U/\mathcal{D}_U \cdot \partial_z$, by sending f to $nfe^{-z^{-n}}$.

Regularity on smooth curves

Let C be a smooth curve. We can then compactify it: $i : C \hookrightarrow \overline{C}$.
(E.g., $\mathbf{A}^1 \hookrightarrow \mathbf{P}^1 = \mathbf{A}^1 \cup \{\infty\}$.)

Let $D = \overline{C} \setminus C$, and let $\mathcal{D}_{\overline{C}}^{Z}$ denote the subsheaf of $\mathcal{D}_{\overline{C}}$ generated by $\mathcal{O}_{\overline{C}}$ and vector fields vanishing on Z.

Definition

A $\mathcal{O}_C$-coherent $\mathcal{D}_C$-module $\mathcal{F}$ is said to have *regular singularities* if $\int_i \mathcal{F}$ is a union of $\mathcal{O}_{\overline{C}}$-coherent $\mathcal{D}_{\overline{C}}^{Z}$-submodules.

A theorem of Deligne's says that having regular singularities is independent of the choice of compactification.

Definition

A holonomic $\mathcal{D}_C$-module $\mathcal{F}$ is said to have regular singularities if there exists an open dense $U \subseteq C$ such that $\mathcal{F}|_U$ is a $\mathcal{O}_U$-coherent $\mathcal{D}_U$-module[a] which has regular singularities.

[a] Remember: this exists!

Little drops of water make a mighty ocean

The definition of regularity in higher dimensions will basically be built from curves. Recall:

Theorem

Let X be a smooth algebraic variety. Then:

1. Let $Y \subseteq X$ be a locally closed smooth connected subvariety of X such that the inclusion $i : Y \hookrightarrow X$ is affine. Let $\mathcal{F}$ be a simple holonomic $\mathcal{D}_X$-module. Then the unique simple submodule of $\int_i \mathcal{F}$ is
$$\mathcal{L}(Y, \mathcal{F}) = i_{!*}\mathcal{F} = \mathrm{im}\left(\int_{i_!} \mathcal{F} \to \int_i \mathcal{F}\right),$$
called the minimal/Goresky-MacPherson extension.
2. Any simple holonomic $\mathcal{D}_X$-module is isomorphic to the minimal extension $i_{!*}\mathcal{F}$, where $i : Y \hookrightarrow X$ is as above, and $\mathcal{F}$ is a simple $\mathcal{O}_Y$-coherent $\mathcal{D}_Y$-module.

Let's get groovy

Let X be a smooth **C**-variety.

1. A $\mathcal{O}_X$-coherent $\mathcal{D}_X$-module $\mathcal{F}$ is said to be regular holonomic if for any smooth curve $C \hookrightarrow X$, the restriction $\mathcal{F}|_C$ is regular holonomic.
2. A simple holonomic $\mathcal{D}_X$-module $\mathcal{F}$ is said to be regular holonomic if it is the minimal extension $i_{!*}(\mathcal{G})$ of an embedding $i : Y \hookrightarrow X$ of a closed smooth subvariety with $\mathcal{G}$ a regular holonomic $\mathcal{D}_Y$-module.
3. A holonomic $\mathcal{D}_X$-module $\mathcal{F}$ is said to be regular holonomic if every simple subquotient of $\mathcal{F}$ is regular holonomic.
4. An object $\mathcal{F} \in \mathrm{D}^b(\mathcal{D}_X)$ is said to be regular holonomic if each cohomology sheaf is regular holonomic.

We'll write $\mathrm{D}^b_{\mathrm{rhol}}(\mathcal{D}_X)$ to denote the full subcategory of $\mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_X)$ spanned by the regular holonomic D-modules.

Examples

We defined regularity so that if $X = \mathbf{A}^1$, and $\mathcal{F} = \mathcal{D}_X/\mathcal{D}_X \cdot P$ (which is holonomic), then $\mathcal{F}$ is regular if and only if the differential equation $Pf = 0$ is regular on $\mathbf{P}^1$ (in the usual sense).

One can think of regularity as imposing a growth condition on solutions. For instance, we saw that $\mathcal{D}_X/\mathcal{D}_X \cdot (z^{n+1}\partial_z - n)$ is not regular.
Let's try to solve the associated differential equation in power series. So we're looking at the kernel of

$$\mathbf{C}[\![z]\!] \xrightarrow{z^{n+1}\partial_z - n} \mathbf{C}[\![z]\!].$$

It's easy to see that there's nothing in the kernel.

Regularity as a growth condition

The cokernel of our map

$$\mathbf{C}[\![z]\!] \xrightarrow{z^{n+1}\partial_z - n} \mathbf{C}[\![z]\!].$$

is interesting: you can solve the recursion defined by

$$\sum_{k\geq 0} b_k z^k = (z^{n+1}\partial_z - n)\sum_{k\geq 0} a_k z^k = \sum_{k\geq n+1}((k-n-1)a_{k-n-1} - nka_k)z^k - \sum_{k=0}^{n} nka_k z^k.$$

But if we tried to solve this in $\mathbf{C}[z]$, we'd fail. E.g., suppose we're solving

$$z = (z^2\partial_z - 1)\sum_{k\geq 0} a_k z^k,$$

then $a_0 = 0$, $a_1 = -1$, and $a_n = (n-1)a_{n-1}$. This means that $a_n = -(n-1)!$, and so

$$\sum_{k\geq 0} a_k z^k = -\sum_{k\geq 0}(k-1)!z^k.$$

The radius of convergence is zero.

Important properties

Theorem

Let X be a smooth algebraic variety. Then $\mathbf{D}$ preserves regular holonomicity, as do direct and exceptional direct image, and inverse and exceptional inverse image.

Theorem

A holonomic $\mathcal{D}_X$-module $\mathcal{F}$ is regular if and only if $i_C^\dagger \mathcal{F}$ is a regular holonomic $\mathcal{D}_C$-module for every locally closed embedding $i : C \hookrightarrow X$ of a smooth curve X.

Some people take this as the definition of regularity.

So you should think of a regular holonomic as a holonomic $\mathcal{D}_X$-module whose restriction to each curve is generically a vector bundle with integrable connection, whose worst singularities are simple poles (i.e., of order 1).

Another important result

Here's another important theorem of (you guessed it) Deligne's.

Theorem (Deligne)

Let X be a smooth variety over $\mathbf{C}$. Then analytification defines an equivalence

$$\left\{\begin{matrix}\text{Vector bundles on } X + \\ \text{flat/integrable connection} \\ \text{which is regular}\end{matrix}\right\} \xrightarrow{\simeq} \left\{\begin{matrix}\text{Vector bundles on } X^{\mathrm{an}} + \\ \text{flat/integrable connection}\end{matrix}\right\}.$$

As a consequence, we get:

Corollary

There's an equivalence

$$\left\{\begin{matrix}\text{Vector bundles on } X + \\ \text{flat/integrable connection} \\ \text{which is regular}\end{matrix}\right\} \xrightarrow{\simeq} \left\{\text{Local systems on } X^{\mathrm{an}}\right\}.$$

Remarks on Deligne's theorem

I won't say much about this, because I haven't read the proof. Roughly, the proof goes as follows:

- Fix a compactification $X \hookrightarrow \overline{X}$ such that $\overline{X} \setminus X$ has simple normal crossings. (Can get this by applying resolution of singularities to a projective compactification of X.)
- Show the key lemma: for any $(\mathcal{F}^{\mathrm{an}}, \nabla^{\mathrm{an}})$ on X^{an}, there is a coherent *analytic* regular $(\overline{\mathcal{F}^{\mathrm{an}}}, \overline{\nabla^{\mathrm{an}}})$ on $\overline{X}^{\mathrm{an}}$ extending $(\mathcal{F}^{\mathrm{an}}, \nabla^{\mathrm{an}})$;
- Apply GAGA to $(\overline{\mathcal{F}^{\mathrm{an}}}, \overline{\nabla^{\mathrm{an}}})$ on $\overline{X}^{\mathrm{an}}$ to get a regular $(\overline{\mathcal{F}}, \nabla)$ on $\overline{X}$;
- Restrict to X.

Let's press pause on regular D-modules for now, and turn to constructibility.

Motivation from D-modules

Recall that holonomic $\mathcal{D}_X$-modules $\mathcal{F}$ are precisely those for which there exist a stratification over which $\mathcal{F}$ corresponds to a flat connection on each stratum.

So, we decided, a "solution sheaf" that might be associated to $\mathcal{F}$ should be such that there exists a stratification on which the solution sheaf is locally constant.

To see this in action, let's look at an example.

Example

Consider the affine curve $X = \mathbf{A}^1_{\mathbf{C}}$ (so $X^{\mathrm{an}} = \mathbf{C}$), and look at the $\mathcal{D}_X$-module $\mathcal{D}_X/\mathcal{D}_X \cdot (z\partial_z - \lambda)$. (We saw that this was a regular holonomic $\mathcal{D}_X$-module.)
The associated differential equation is

$$z\partial_z f = \lambda f. \tag{2}$$

We'll look at the sheaf $\mathcal{P}$ of solutions to (2).

Let $j : \mathbf{C}^\times \hookrightarrow \mathbf{C}$ denote the inclusion.
The solutions to the differential equation (2) are cz^λ, with $c \in \mathbf{C}$.
It follows that $\mathcal{P}|_{\mathbf{C}^\times}$ forms a rank one local system.

If $\lambda \notin \mathbf{Z}_{\geq 0}$, then there is no solution to the differential equation $z\partial_z f = \lambda f$ on $\mathbf{C}$, so $\mathcal{P}|_{\{0\}} = 0$.

But if $\lambda \in \mathbf{Z}_{\geq 0}$, then we can in fact solve the differential equation $z\partial_z f = \lambda f$ on $\mathbf{C}$, by cz^λ with $c \in \mathbf{C}$. So:

$$\mathcal{P}|_{\{0\}} \cong \begin{cases} \mathbf{C} & \lambda \in \mathbf{Z}_{\geq 0} \\ 0 & \text{else.} \end{cases}$$

Defining constructibility

Note that if λ is not an integer, then z^λ is not an algebraic function — and so we really need to work in the analytic topology to see the phenomenon described above.

What this example showed us was that a differential equation on X naturally gives a "solution sheaf" $\mathcal{P}$ on X^{an} for which there exist a stratification, such that $\mathcal{P}$ is locally constant on each stratum.
This'll be the definition of constructibility.

Stratification

Let X be an algebraic variety. A locally finite partition $X = \coprod_{\alpha \in A} X_\alpha$ by locally closed subvarieties is called a stratification if each X_α is smooth, and $\overline{X}_\alpha = \coprod_{\beta \in B} X_\beta$ for some $B \subseteq A$.

Constructibility

Let X be an algebraic variety. A **C**-module $\mathcal{P}$ on X^{an} is said to be constructible if there exists a stratification $\coprod_\alpha X_\alpha$ such that each $\mathcal{F}|_{X_\alpha^{\mathrm{an}}}$ is locally constant.

Examples

Obviously, all local systems give examples of constructible sheaves.
Fun example: let $\mathcal{L}$ denote the local system associated to the representation

$$\pi_1(\mathbf{C} \setminus \{0\}) = \mathbf{Z} \to \mathbf{C}^\times,\ 1 \mapsto -1.$$

This is the "square root sheaf": if $f : \mathbf{C} \setminus \{0\} \to \mathbf{C} \setminus \{0\}$ is the degree two covering, then

$$f_*(\mathbf{C}) \cong \mathbf{C} \oplus \mathcal{L}.$$

The degree two covering extends to $\mathbf{C}$, except with ramification at zero: in other words, we still have the map $f : \mathbf{C} \to \mathbf{C}$ sending z to z^2.
Then, if $X \subseteq \mathbf{C}$, we have:

$$f_*(\mathbf{C})|_X \cong \begin{cases} \mathbf{C} \oplus \mathcal{L} & X = \mathbf{C} \setminus \{0\} \\ \mathbf{C} & X = \{0\}. \end{cases}$$

So $f_*(\mathbf{C})$ is a constructible sheaf on $\mathbf{C}$.

In fact, constructible sheaves on a curve C are just specified by a finite set of points $\{x_i\}$, a local system on $C \setminus \{x_i\}$, a vector space V_i over each x_i, and "gluing data".

Sometimes you can be lazy and correct

As usual, we say that a complex $\mathcal{F} \in \mathrm{D}^b(X)$ is constructible if all its cohomology sheaves are constructible.
One nice fact is that the bounded derived categories of:

- complexes with constructible cohomology;
- constructible sheaves

are actually equivalent. So you can just say "bounded derived category of constructible sheaves" without any ambiguity.
We'll denote this category by $\mathrm{D}^b_c(X)$. Whenever we say constructible sheaves, we'll *always* be derived.

We'd like to define functors on constructible sheaves, like in the previous talk.

This is where we stopped on 3/21.

Recap of yesterday

In case you might've forgotten, here's a quick summary of what happened yesterday.

Riemann-Hilbert

Let X be a complex manifold. Then there is an equivalence of categories:

$$\{\text{Local systems on } X\} \xrightarrow{\sim} \left\{\begin{matrix}\text{Vector bundles on } X + \\ \text{flat/integrable connection}\end{matrix}\right\};$$

$$\mathcal{L} \mapsto (\mathcal{L} \otimes_{\mathbf{C}} \mathcal{O}_X, 1 \otimes d);$$

$$\ker(\nabla) \leftarrow\!\shortmid (\mathcal{F}, \nabla)$$

We then determined you needed to impose regularity and holonomicity conditions to generalize to the algebraic setting, and to allow for more D-modules.

Recap, continued

Regarding holonomicity, we stated the following theorem, which, in my opinion, is the most conceptual way of thinking about holonomicity.

Theorem

Let $\mathcal{F} \in \mathrm{D}^b(\mathcal{D}_X)$. TFAE:

- $\mathcal{F}$ is holonomic;
- for each $i_x : \{x\} \hookrightarrow X$, the complex $i_x^\dagger \mathcal{F}$ is finite-dimensional;
- there is a finite sequence

$$\emptyset = X_{n+1} \subseteq X_n \subseteq \cdots \subseteq X_0 = X$$

  of closed subsets with $X_r \setminus X_{r+1}$ smooth, and each $i_r^\dagger \mathcal{F}$ a coherent $\mathcal{O}_X$-module, where $i_r : X_r \setminus X_{r+1} \hookrightarrow X$.

More recapping

We defined what it means for a holonomic D-module to be *regular*, first by defining it for curves, and then defining it for all smooth varieties.

Rather than give you the definition again, let me recall some examples.

Over $X = \mathbf{A}^1$, a $\mathcal{D}_X$-module $\mathcal{D}_X/\mathcal{D}_X \cdot P$ with P defining a differential equation

$$\frac{d^n f}{dz^n} + p_1(z)\frac{d^{n-1} f}{dz^{n-1}} + \cdots + p_n(z) f = 0,$$

with each p_i meromorphic.

Definition

This differential equation has a *regular singularity* at z_0 if:

- not all p_k are analytic at z_0;
- all $(z - z_0)^k p_k(z)$ are analytic at z_0, i.e., p_k has a pole of order at most k at z_0.

A bit more recapping

For instance, the following $\mathcal{D}_X$-modules are regular (over $X = \mathbf{A}^1$):

- $$\mathcal{F} = \mathcal{D}_X/\mathcal{D}_X \cdot (z\partial_z - \lambda) = \mathcal{O}_X\{z^\lambda\}.$$
- $$\mathcal{F} = \mathcal{O}_U\{\log(z)\}.$$

However,

$$\mathcal{F} = \mathcal{D}_U/\mathcal{D}_U \cdot (z^{n+1}\partial_z + n) = \mathcal{O}_U\{\exp(z^{-n})\}$$

is not regular.

We also said that regularity imposed a growth condition on the solutions to our differential equation, if that's your vibe.

Almost done

Finally, by looking at the differential equation $z\partial_z f = \lambda f$ over $\mathbf{A}^1$, we saw that the sheaf $\mathcal{P}$ of solutions was isomorphic to the constant sheaf $\mathbf{C}$ over $\mathbf{A}^1 \setminus \{0\}$, but was either 0 over $\{0\}$ (if $\lambda \notin \mathbf{Z}_{\geq 0}$) or $\mathbf{C}$ (if $\lambda \in \mathbf{Z}_{\geq 0}$).

This led to the notion of constructibility.

Constructibility

Let X be an algebraic variety. A $\mathbf{C}$-module $\mathcal{P}$ on X^{an} is said to be constructible if there exists a stratification $\coprod_\alpha X_\alpha$ such that each $\mathcal{F}|_{X_\alpha^{\mathrm{an}}}$ is locally constant.

Another example to keep in mind is: the pushforward of the constant sheaf $\mathbf{C}$ along the degree two map $\mathbf{C} \setminus \{0\} \to \mathbf{C} \setminus \{0\}$.

Now we're going to move on, and define the six functor formalism on constructible sheaves, and then state the full Riemann-Hilbert correspondence.

Four functors

Let $f : X \to Y$ be a morphism of finite type. It's easy to see that $f^* : \mathrm{D}^b(Y) \to \mathrm{D}^b(X)$ preserves constructibility, as does f_* — but this latter fact is hard (essentially boils down to Chevalley constructibility).

Suppose f is separated. Nagata's compactification theorem tells us that there is a proper morphism $\overline{f} : \overline{X} \to Y$ and an open immersion $j : X \hookrightarrow \overline{X}$ such that

$$f : X \hookrightarrow \overline{X} \xrightarrow{\overline{f}} Y.$$

Given $\mathcal{F} \in \mathrm{D}^b(X)$, define

$$f_!(\mathcal{F}) = \overline{f}_* j_!(\mathcal{F}),$$

where $j_!$ is extension by zero. One can show that $f_!$ preserves constructibility.

One can also show that $f_!$ admits a right adjoint $f^!$ (and this also preserves constructibility).

$4 + 2 = 6$

When $f : X \to *$, we get the dualizing complex $f^! \overline{\mathbf{C}} \in \mathrm{D}^b_c(X)$. Define $\mathbf{D}_X : \mathrm{D}^b(X) \to \mathrm{D}^b(X)^{\mathrm{op}}$ via

$$\mathbf{D}_X(\mathcal{F}) = \mathrm{Hom}_{\mathbf{C}}(\mathcal{F}, f^! \overline{\mathbf{C}}).$$

Again, this preserves constructibility.

Finally, if X and Y are analytic varieties, and $\mathcal{F} \in \mathrm{D}^b(X)$ and $\mathcal{G} \in \mathrm{D}^b(Y)$, we have

$$\mathcal{F} \boxtimes_{\mathbf{C}} \mathcal{G} = p_1^{-1}\mathcal{F} \otimes_{\mathbf{C}} p_2^{-1}\mathcal{G} \in \mathrm{D}^b(X \times Y),$$

where $p_1 : X \times Y \to X$ and $p_2 : X \times Y \to Y$.

The sixfold way

In summary, if $f : X \to Y$ is a separated morphism of finite type, we get

$$f_*, f_! : \mathrm{D}^b_c(X) \to \mathrm{D}^b_c(Y),\ f^*, f^! : \mathrm{D}^b_c(Y) \to \mathrm{D}^b_c(X),$$
$$\mathbf{D} : \mathrm{D}^b_c(X) \xrightarrow{\sim} \mathrm{D}^b_c(X)^{\mathrm{op}},\ -\boxtimes- : \mathrm{D}^b_c(X) \times \mathrm{D}^b_c(Z) \to \mathrm{D}^b_c(X \times Z).$$

Solving differential equations

There's a lot more to say, but we have to move on lest we don't get to our destination.

We keep emphasizing that the "solution sheaf" associated to a holonomic D-module is constructible — that's how we *motivated* the definition of constructibility. We should state that precisely.
First, we need to define the solution sheaf. Suppose $\mathcal{F} = \mathcal{D}_X/\mathcal{D}_X \cdot P$; then, there is an exact sequence

$$\mathsf{Hom}_{\mathcal{D}_X}(\mathcal{F}, \mathcal{O}_X) \to \mathsf{Hom}_{\mathcal{D}_X}(\mathcal{D}_X, \mathcal{O}_X) \cong \mathcal{O}_X \xrightarrow{P} \mathcal{O}_X \cong \mathsf{Hom}_{\mathcal{D}_X}(\mathcal{D}_X, \mathcal{O}_X),$$

so $\mathsf{Hom}_{\mathcal{D}_X}(\mathcal{F}, \mathcal{O}_X)$ is $\{f \in \mathcal{O}_X | Pf = 0\}$.

Definition

Let X be a smooth $\mathbf{C}$-variety, and let $\mathcal{F}$ be a $\mathcal{D}_X$-module. Define

$$\mathrm{Sol}(\mathcal{F}) = \mathsf{Hom}_{\mathcal{D}_{X^{\mathrm{an}}}}(\mathcal{F}^{\mathrm{an}}, \mathcal{O}_{X^{\mathrm{an}}}),$$

where we're taking derived Hom.

Kashiwara's constructibility theorem

Here's the all-important result.

Theorem (Kashiwara)

Let X be a smooth $\mathbf{C}$-variety, and let $\mathcal{F}$ be a holonomic $\mathcal{D}_X$-module; then $\mathrm{Sol}(\mathcal{F})$ is a constructible sheaf.

We will not describe the proof. But we'll see examples of this in action below.

Riemann-Hilbert, redux

Kashiwara's theorem tells us that Sol defines a functor $\mathrm{D}^b_{\mathrm{hol}}(\mathcal{D}_X) \to \mathrm{D}^b_c(X)^{\mathrm{op}}$ from holonomic $\mathcal{D}_X$-modules to constructible sheaves. It's not an equivalence on the entire category, but:

The Riemann-Hilbert correspondence

Let X be a smooth $\mathbf{C}$-variety. Then the functor

$$\mathrm{Sol} : \mathrm{D}^b_{\mathrm{rhol}}(\mathcal{D}_X) \to \mathrm{D}^b_c(X)^{\mathrm{op}}$$

is an equivalence of derived categories.

This the contravariant Riemann-Hilbert correspondence. There's a covariant version, where one composes with the duality $\mathbf{D}$.

The de Rham functor

More precisely, there is an isomorphism

$$\mathrm{Sol}(\mathbf{D}(\mathcal{F})) = \mathsf{Hom}_{\mathcal{D}_X}(\mathbf{D}(\mathcal{F}), \mathcal{O}_X) \cong \mathsf{Hom}_{\mathcal{D}_X}(\mathcal{O}_X, \mathcal{F}) \cong \mathsf{Hom}_{\mathcal{D}_X}(\mathcal{O}_X, \mathcal{D}_X) \otimes_{\mathcal{D}_X} \mathcal{F},$$

so we need to resolve $\mathcal{O}_X$ as a $\mathcal{D}_X$-module.

To do this, we have to resolve $\mathcal{O}_X$ as a $\mathcal{D}_X$-module.
If Θ_X denotes the tangent sheaf, then

$$0 \to \mathcal{D}_X \otimes_{\mathcal{O}_X} \bigwedge^n \Theta_X \to \cdots \to \mathcal{D}_X \otimes_{\mathcal{O}_X} \bigwedge^0 \Theta_X \cong \mathcal{D}_X \to \mathcal{O}_X \to 0$$

is a resolution of $\mathcal{O}_X$ as a $\mathcal{D}_X$-module.

So,

$$\mathsf{Hom}_{\mathcal{D}_X}(\mathcal{O}_X, \mathcal{D}_X) \cong \left[\mathsf{Hom}_{\mathcal{O}_X}\left(\bigwedge^0 \Theta_X, \mathcal{D}_X\right) \to \cdots \to \mathsf{Hom}_{\mathcal{O}_X}\left(\bigwedge^n \Theta_X, \mathcal{D}_X\right)\right].$$

The de Rham functor, continued

This is in turn isomorphic to the complex

$$\left[\Omega^0_X \otimes_{\mathcal{O}_X} \mathcal{D}_X \cong \mathcal{D}_X \to \cdots \to \Omega^n_X \otimes_{\mathcal{O}_X} \mathcal{D}_X\right],$$

which, it turns out, is a resolution of $\Omega^n_X[-n] = \omega_X$ as a *right* $\mathcal{D}_X$-module.

The upshot is that $\mathrm{Hom}_{\mathcal{D}_X}(\mathcal{O}_X, \mathcal{D}_X) \cong \omega_X$ in the derived category, and so

$$\mathrm{Sol}(\mathbf{D}(\mathcal{F})) \cong \mathrm{Hom}_{\mathcal{D}_X}(\mathcal{O}_X, \mathcal{D}_X) \otimes_{\mathcal{D}_X} \mathcal{F} \cong \omega_X \otimes_{\mathcal{D}_X} \mathcal{F}.$$

de Rham complex

Let X be a smooth $\mathbf{C}$-variety, and let $\mathcal{F}$ be a holonomic $\mathcal{D}_X$-module. The de Rham complex $\mathrm{DR}(\mathcal{F})$ is

$$\mathrm{DR}(\mathcal{F}) = \omega_{X^{\mathrm{an}}} \otimes_{\mathcal{D}_{X^{\mathrm{an}}}} \mathcal{F}^{\mathrm{an}}.$$

Kashiwara implies that $\mathrm{DR}(\mathcal{F})$ is constructible, and the Riemann-Hilbert correspondence equivalently states that $\mathrm{DR} : \mathrm{D}^b_{\mathrm{rhol}}(\mathcal{D}_X) \to \mathrm{D}^b_c(X)$ is an equivalence.

An example

Let's return to our favorite example, $X = \mathbf{A}^1$ (so $X^{\mathrm{an}} = \mathbf{C}$). Let $\mathcal{F} = \mathcal{D}_X/\mathcal{D}_X \cdot (z\partial_z - \lambda)$.
In motivating the definition of constructibility, we saw that if $\mathcal{P} = \mathrm{H}^0(\mathrm{Sol}(\mathcal{F}))$, then $\mathcal{P}|_{\mathbf{C}\setminus\{0\}} = \mathbf{C}$, and

$$\mathcal{P}|_{\{0\}} = \begin{cases} \mathbf{C} & \lambda \in \mathbf{Z}_{\geq 0} \\ 0 & \text{else.} \end{cases}$$

The actual complex $\mathrm{Sol}(\mathcal{F})$ is

$$\mathbf{C}[z] \xrightarrow{z\partial_z - \lambda} \mathbf{C}[z].$$

If, instead, $\mathcal{F} = \mathcal{D}_X/\mathcal{D}_X \cdot \partial_z$, then

$$\mathrm{Sol}(\mathcal{F}) = [\mathbf{C}[z] \xrightarrow{\partial_z} \mathbf{C}[z]],$$

and this only has cohomology in degree 0 (where it's $\mathbf{C}$).

An example, continued

Finally, suppose $\mathcal{F} = \mathcal{D}_X/\mathcal{D}_X \cdot z$. Then

$$\mathrm{Sol}(\mathcal{F}) = [\mathbf{C}[z] \xrightarrow{z} \mathbf{C}[z]],$$

which has cohomology in degree 1 and is supported at $z = 0$.

Observation

For all these cases, only one cohomology sheaf is nonzero. If it's in degree 1, then it's supported at $z = 0$; if it's in degree 0, it's supported on a 1-dimensional piece.

This is an example of Kashiwara's constructibility theorem in action.

Proving the Riemann-Hilbert correspondence

The hard input into the Riemann-Hilbert correspondence is the following theorem, which we shall not prove.

A hard theorem

Let $f : X \to Y$ be a morphism of smooth algebraic varieties. The de Rham complex functor $\mathrm{DR} : \mathrm{D}^b_{\mathrm{rhol}}(\mathcal{D}_X) \to \mathrm{D}^b_c(X^{\mathrm{an}})$ commutes with duals, direct and exceptional direct image, inverse image and exceptional inverse image, and products.

Here, it is critical that one uses regular $\mathcal{D}_X$-modules.

Now we're cruising: essential surjectivity

The essential surjectivity of DR is rather easy. It suffices to check that the generators of $\mathrm{D}^b_c(X^{\mathrm{an}})$.

We claim that $\mathrm{D}^b_c(X^{\mathrm{an}})$ is generated by $i_*\mathcal{L}$ for a closed embedding $Z \hookrightarrow X$ of a locally closed smooth subvariety and a local system $\mathcal{L}$ on Z^{an}.

(To see this, let $\mathcal{F} \in \mathrm{D}^b_c(X^{\mathrm{an}})$, and let $j : \mathsf{Supp}(\mathcal{F}) \hookrightarrow X$. Then the fiber of $\mathcal{F} \to j_*j^*\mathcal{F}$ has support strictly smaller than $\mathsf{Supp}(\mathcal{F})$ — now induct!)

So it suffices to show that there is a $\mathcal{D}_Z$-module $\mathcal{F}$ on Z such that $\mathrm{DR}(\mathcal{F}) = \mathcal{L}$. But this is Deligne's Riemann-Hilbert correspondence (along with the duality between Sol and DR).

Full faithfulness given the hard theorem

We just calculate. Let $\Delta : X \to X \times X$ denote the diagonal embedding, and let $p : X \to *$ denote the projection to a point.

If $\mathcal{L}, \mathcal{L}' \in \mathrm{D}^b_c(X^{\mathrm{an}})$, then:

$$\begin{aligned} p_* \operatorname{Hom}(\mathcal{L}, \mathcal{L}') &\simeq p_* \operatorname{Hom}(\mathcal{L}, \mathbf{D}^2(\mathcal{L}')) \\ &\simeq p_* \operatorname{Hom}(\mathcal{L} \otimes \mathbf{D}(\mathcal{L}'), \omega_X) \\ &\simeq p_* \mathbf{D}(\mathcal{L} \otimes \mathbf{D}(\mathcal{L}')) \\ &\simeq p_* \mathbf{D} \Delta^{-1}(\mathcal{F} \boxtimes \mathbf{D}(\mathcal{L}')) \\ &\simeq p_* \Delta^!(\mathbf{D}(\mathcal{F}) \boxtimes \mathcal{L}'). \end{aligned}$$

The same string of identifications (in the setting of D-modules) shows that if $\mathcal{F}, \mathcal{G} \in \mathrm{D}^b_{\mathrm{rhol}}(X)$, then

$$\int_p \Delta^!(\mathbf{D}_X(\mathcal{F}) \boxtimes \mathcal{G}) \simeq \operatorname{Hom}_{\mathcal{D}_X}(\mathcal{F}, \mathcal{G}).$$

We're done!

It follows that if $\mathcal{F}, \mathcal{G} \in \mathrm{D}^b_{\mathrm{rhol}}(X)$, then (because DR commutes with everything by the hard theorem):

$$\begin{aligned}\mathrm{Hom}(\mathrm{DR}(\mathcal{F}), \mathrm{DR}(\mathcal{G})) &\simeq p_* \Delta^! (\mathbf{D}_X(\mathrm{DR}(\mathcal{F})) \boxtimes \mathrm{DR}(\mathcal{G})) \\ &\simeq \mathrm{DR}_{\mathrm{pt}} \int_p \Delta^! (\mathbf{D}_X(\mathcal{F}) \boxtimes \mathcal{G}) \\ &\simeq \int_p \Delta^! (\mathbf{D}_X(\mathcal{F}) \boxtimes \mathcal{G}) \simeq \mathrm{Hom}_{\mathcal{D}_X}(\mathcal{F}, \mathcal{G}).\end{aligned}$$

Hearts and lungs

The derived category of $\mathcal{D}_X$-modules admits a *t*-structure, and this descends to $\mathrm{D}^b_{\mathrm{rhol}}(\mathcal{D}_X)$.

If the word "*t*-structure" is unfamiliar, think of it as a division of the derived category into complexes with cohomology concentrated in positive and negative degrees. So if $\mathcal{A}$ is an abelian category, then $\mathrm{D}^+(\mathcal{A})$ and $\mathrm{D}^-(\mathcal{A})$ defines a t-structure on $\mathrm{D}(\mathcal{A})$.
The heart of the *t*-structure is $\mathrm{D}^+(\mathcal{A}) \cap \mathrm{D}^-(\mathcal{A}) \simeq \mathcal{A}$.

In particular, the heart of the *t*-structure on $\mathrm{D}^b_{\mathrm{rhol}}(\mathcal{D}_X)$ is the abelian category of regular holonomic $\mathcal{D}_X$-modules.

We can transport this *t*-structure to $\mathrm{D}^b_c(X)$ via DR — so what is it? What's the heart? It's *not* just sheaves of $\mathbf{C}$-modules on X concentrated in degree zero.

Perverse sheaves

Definition

Let $\mathcal{P} \in \mathrm{D}^b_c(X)$. Say that $\mathcal{P}$ is a perverse sheaf if there exists $\mathcal{F} \in \mathrm{D}^b_{\mathrm{rhol}}(\mathcal{D}_X)$ such that $\mathrm{DR}(\mathcal{F}) \cong \mathcal{P}$.

Obviously, this is an unsatisfactory definition: we'd like a definition intrinsic to constructible sheaves.
Recall our observation from a few slides back:

Observation

Over $\mathbf{A}^1$, only one cohomology sheaf of $\mathrm{Sol}(\mathcal{F})$ for "easy" $\mathcal{F}$ turned out to be nonzero. If it's concentrated in degree 1, then it's supported at $z = 0$; if it's concentrated in degree 0, it's supported on a 1-dimensional piece.

This actually turns out to be true over $\mathbf{A}^n$, too.

Affine space

Let $X = \mathbf{A}^n$. Suppose $\mathcal{F} = \mathcal{D}_X/\mathcal{D}_X \cdot (P_1, \cdots, P_n)$ with each P_i of the form $z_i\partial_{z_i} - \lambda$, z_i, or ∂_{z_i}.

By arguing exactly as in the case $n = 1$, one finds that the cohomology sheaves of $\mathrm{Sol}(\mathcal{F})$ are zero everywhere except in dimension k, where k is the number of indices i such that $P_i = z_i$.
As in the case $n = 1$,

$$\mathsf{Supp}(\mathrm{H}^k(\mathrm{Sol}(\mathcal{F}))) = V(z_i | P_i = z_i).$$

Note that this is a codimension k subvariety of $\mathbf{A}^n$.

For any $\mathcal{F}$ built from these "simple" $\mathcal{D}_X$-modules by extensions, an argument with the cohomology long exact sequence shows that

$$\operatorname{codim} \mathsf{Supp}(\mathrm{H}^k(\mathrm{Sol}(\mathcal{F}))) \geq k.$$

(What we're seeing again is Kashiwara's constructibility theorem in action.)

Pervy

In terms of the de Rham functor, we have

$$\dim \mathrm{Supp}(\mathrm{H}^k(\mathrm{DR}(\mathcal{F}))) \leq -k.$$

You can also do this analysis for $\mathbf{D}(\mathrm{DR}(\mathcal{F})) \simeq \mathrm{DR}(\mathbf{D}(\mathcal{F}))$, and this gives

$$\dim \mathrm{Supp}(\mathrm{H}^k(\mathbf{D}(\mathrm{DR}(\mathcal{F})))) \leq -k.$$

It turns out (Kashiwara) that this characterizes regular holonomic $\mathcal{D}_X$-modules, thought of as sitting inside $\mathrm{D}^b_{\mathrm{rhol}}(\mathcal{D}_X)$.

So:

Definition/Theorem

Let $\mathcal{P} \in \mathrm{D}^b_c(X)$. Then $\mathcal{P}$ is a perverse sheaf iff

$$\dim \mathrm{Supp}(\mathrm{H}^k(\mathcal{F})) \leq -k,$$
$$\dim \mathrm{Supp}(\mathrm{H}^k(\mathbf{D}(\mathcal{F}))) \leq -k.$$

I like how the nlab describes perverse sheaves: "They are neither perverse nor sheaves."

Also: "in some languages ... such as German, it sounds no better than 'idiotic sheaf' or the like".

And Grothendieck himself said:

> *What an idea to give such a name to a mathematical thing! Or to any other thing or living being, except in sternness towards a person — for it is evident that of all the 'things' in the universe, we humans are the only ones to whom this term could ever apply.*

The last line surely says something about Grothendieck's private life...

Some sources

- Most comprehensive source: Hotta-Takeuchi-Tanisaki's *D-Modules, Perverse Sheaves, and Representation Theory*.
- Notes from a course by Christian Schnell, `http://www.math.stonybrook.edu/~cschnell/mat615/`. These are really nice; I learnt a lot of examples from here.
- Notes from a course by Sergey Arkhipov, `http://www.unige.ch/math/folks/nikolaev/assets/files/130726232840.pdf`.
- Katz, *An Overview of Deligne's Work on Hilbert's Twenty-First Problem*. Available at `https://web.math.princeton.edu/~nmk/old/DeligneXXIHilbert.pdf`.

I also wrote up some notes at `http://www.mit.edu/~sanathd/riemann-hilbert.pdf`, but they don't include background on constructible sheaves, or details about perverse sheaves.