GROTHENDIECK AND SPRINGER'S NEW YEAR RESOLUTIONS FOR 2019

ROK GREGORIC

This is a LaTeXed version of an email I sent on January 16, 2019 to my friend Tom Gannon, an aspiring geometric group theorist and fellow grad student at UT. I had just spent a week in New York city, visiting my coadvisor at Columbia, and learning cool, albeit rather basic, representation theory. Without anybody there to share the excitement, I channelled it into this email to Tom.

The topic of this note is mostly the Springer resolution, and its Grothendieckian upgrade, although there is a rather lengthy digression into the most basic basics of Hamiltonian mechanics in there for good measure. While no knowledge of the latter is presupposed, some basic familiarity (on the level of Fulton-Harris or lower) with representation theory is assumed. In particular, we trust the reader to know what we mean by saying that we fix a Lie group $\mathfrak{g}$, and possibly inside it a Borel subalgebra $\mathfrak{b}$ and Cartan subalgebra $\mathfrak{h}$.

As always, I can not guarantee that everything below is correct, only that I thought it was when I wrote it. Use at your own risk!

1. Nilpotent Co' Ain't Yo Foe

The classical Springer resolution story can be told in many ways, but we have to start somewhere: consider the nilpotent cone

$$N = \{\text{elements } x \text{ in } \mathfrak{g} \text{ which are nilpotent, i.e for which } \mathrm{ad}(x) \text{ is nilpotent}\}.$$

Surely we like N; given a Borel $\mathfrak{b}$, we have $\mathfrak{b} \cap N = \mathfrak{n}$, singling out the nilpotent radical. This is cool because it gives us the recipe to extract $\mathfrak{n}$ out of the Borel without any choices, and so we can get hold of the 'abstract' Cartan (I still think that's a silly name) as the quotient $\mathfrak{b}/\mathfrak{n} = \mathfrak{h}$.

Compare this to how the 'more classical' (= Fulton-Harris) way one might go about things: fix a Cartan subalgebra $\mathfrak{h}$ inside $\mathfrak{b}$, this produces a root system and you can take h to be the direct sum of all the root spaces corresponding to strictly positive roots (while $\mathfrak{b}$ and $\mathfrak{h}$ are sums of root spaces for non-negative and zero roots respectively). That's all fine and good and is the way you want to think of the $\mathfrak{h}$, $\mathfrak{b}$, $\mathfrak{n}$ in terms of the weight yoga. Alas, the theoretical shortcoming is that all the identifications, such as $\mathfrak{b}/\mathfrak{n} = \mathfrak{h}$, seemingly depend on the choices we made. Said differently, it isn't obvious from this approach that there is a distinguished isomorphism $\mathfrak{b}/\mathfrak{n} = \mathfrak{h}$, independent of the choice of $\mathfrak{h}$ inside $\mathfrak{b}$. That would a big hinderance, as having this universality is very useful a number of times in rep theoretic considerations you, and to a lesser extent I, engage in.

Btw, it looks at first glance as if we did something illegal somewhere in the first approach: how can $\mathfrak{n}$ inside $\mathfrak{b}$ be uniquely determined without specifying $\mathfrak{h}$? After all, $\mathfrak{h}$ constitutes the orthogonal complement to $\mathfrak{n}$ in $\mathfrak{h}$ wrt the Killing form, and by standing assumption that the Lie algebra $\mathfrak{g}$ we're playing around in is semisimple, the KIlling form is non-degenerate. So it looks as though we've found a universal way of locating the Cartan $\mathfrak{h}$ in $\mathfrak{b}$. Alas, we know that a single Borel may contain more than one Cartan (right?). But

Date: February 4, 2020.
University of Texas at Austin.

I think we're good - the point is that, despite the Killing form being a non-degenerate bilinear form, its restriction to subspaces need not be. Indeed, recall Cartan's criterion: in order for Killing to be non-degenerate[1], the Lie algebra must be semisimple. And yes, the Killing form of $\mathfrak{g}$ restricts to the Killing form of $\mathfrak{b}$, but the Borel subalgebra, nice as it is, certainly isn't semisimple - its radical, far from being trivial, equals itself! So when we said that h was the orthogonal complement to $\mathfrak{n}$ in $\mathfrak{b}$ wrt to the Killing form, that didn't make any sense.

Alright, that was a weird digression. Back to the nilpotent cone. N isn't just a set: the coefficients of the characteristic polynomial $p(t) = \det(t - \mathrm{ad}(x))$ of the linear map $\mathrm{ad}(x)$ are polynomial equations in x, and their common vanishing set is precisely N (indeed, a matrix is nilpotent iff all its eigenvalues are 0 - think about the Jordan form!). This exhibits N as a closed subvariety in $\mathfrak{g}$, so it's a geometrically meaningful guy. Wohoo!

Here's another description of N: you may start with the quotient map $\mathfrak{g} \to \mathfrak{g}//G$ and ask yourself what its (scheme theoretic) fiber over 0 is (spoilers: it's N). Passing to functions, we're asking for the pushout of the diagram $\mathbf{C} \leftarrow \mathrm{Sym}^*(\mathfrak{g})^G \to \mathrm{Sym}^*(\mathfrak{g})$, where the wrong-way map is the augmentation of the symmetric algebra, and the right-way map is the inclusion. The pushout thus is the quotient of $\mathrm{Sym}^*(\mathfrak{g})$ by the ideal generated by $\mathrm{Sym}^+(\mathfrak{g})^G$. Thinking a little about the ideal we are cutting out, and remembering the yoga of symmetric polynomials (that's precisely what we're cutting out), we observe that we are once again cutting out the coefficients of the characteristic polynomial, which we already know to produce N.

How to think about this fact, this fiber presentation of N? Well, it says that N, as a subset of $\mathfrak{g}$, has a unique dense open Ad-orbit, a well-known important fact about the nilpotent cone (in fact, it has finitely many adjoint orbits overall). But the way to think about it: if you have an $x \in \mathfrak{g}$ such that $\mathrm{Ad}(s)(x) = 0$ for all $s \in G$, plug in $s = \exp(y)$ for any $y \in \mathfrak{g}$ to see that $\mathrm{ad}(x) = 0$. But really, you are trying to take the fiber fro the AG point of view, i.e. to evaluate anything, we must first reduce it. That is to say, it would be just as good if instead $\mathrm{ad}(x)^k = 0$ for some $k >> 0$. And lo and behold, that's nothing but nilpotence of x.

Okay, time to move on to the Springer goodness.

2. Springer is German for Knight in Chess

So say we care about N, and know it's a variety. Perhaps sadly, it turns out to have singularities. For elements of $\mathfrak{g} = \mathfrak{sl}_2$ for example, nilpotence = non-maximal rank = vanishing of determinant. Hence N is cut out from $\mathbf{A}^3$ by $x^2 + yz = 0$, so it is the usual cylindrical cone (substitute $y \mapsto y - z$ and $z \mapsto z - y$, if that makes you feel better). In fact the nilpotent cone of any g will be a cone in the sense of algebraic geometry, i.e. a $\mathbf{G}_m$-invariant subvariety of $\mathfrak{g}$, and it will contain the origin 0 in $\mathfrak{g}$, which will hence have to be a singular point.

The Springer resolution $\widetilde{N}$ is, as the name suggests, a resolution, namely a (specific) resolution of singularities for N.

But you don't need to know what you're doing to stumble upon $\widetilde{N}$. If you like the flag variety's propery that $G/B = \{$Borel subalgebras in $\mathfrak{g}\}$, then you might form the following by accident:

$$\widetilde{N} = \{(x, \mathfrak{b}); \mathfrak{b} \text{ a Borel in } \mathfrak{g},\ x \text{ a point in } \mathfrak{n} \text{ (corresponding to the Borel } \mathfrak{b})\}.$$

[1]Math makes us utter some truly shocking strings of words sometimes. The standard joke about birational geometers at an airport comes to mind.

See how useful it was here that we know we can take $\mathfrak{n}$ inside $\mathfrak{b}$ without making any choices! Anyway, clearly $\widetilde{N}$, being an incidence subset of $G/B\mathfrak{g}$, has a map to G/B given by $(x, \mathfrak{b}) \mapsto \mathfrak{b}$, and a map to $\mathfrak{g}$ given by $(x, \mathfrak{b}) \mapsto x$. Note though that, since $\mathfrak{n}$ is always nilpotent, and any nilpotent element belongs to some Borel, the map $\widetilde{N} \to \mathfrak{g}$ factors surjectively through the inclusion $N \to \mathfrak{g}$. Voila, we have obtained a map $\widetilde{N} \to N$, which we claim is the desired resolution of singularities.

The proof sees you restrict to regular elements of N, i.e. such $x \in N$ that $\mathfrak{h} = C_{\mathfrak{g}}(x)$ is a Cartan subalgebra. That's an acceptable thing to do, seeing how the regular elements are dense. For a fixed x in N_{reg}, look at the fiber above it in $\widetilde{N}$ and recognize it's precisely the Borels containing x. But a regular nilpotent element is apparently contained in a unique Borel, so the map $\widetilde{N} \to N$ is indeed a bijection. This has been a very brief proof sketch, and some more comments on what goes into a real proof will be provided later in §7.

Of course, $\widetilde{N}$ is much cooler than just being a resolution of singularities. This can be seen by the sheer myriad of different ways you could stumble on it.

3. You Could Also be a Differential Geometer

Say you were thinking about your friend the flag variety $X = G/B$ as a smooth projective variety, in particular as a smooth manifold. The whole deal of smooth manifolds is that they are geometric objects, by which I mean the underlying topological manifold, but they come equipped with enough extra structure that linear methods, the one thing in mathematics we actually know decently well, are available to study their structure, at least on very very small pieces. (In fact, this same paradigm explains the ambiquity of differential calculus in all its incarnations). The linearization mechanism comes packaged as the (co)tangent bundle. Thus the main thing that viewing X as a smooth manifold instantly buys you, is giving you access to its cotangent bundle T^*X.

What is its fiber of $T^*X \to X$ over a point $\mathfrak{b}$ of X (which we are identifying with the choice of a Borel subalgebra $\mathfrak{b}$ in $\mathfrak{g}$)? Since G and B are both framed manifolds, the sought cotangent space is the quotient $\mathfrak{g}^*/\mathfrak{b}^* = (\mathfrak{n}^-)^*$, where $\mathfrak{n}^-$ is the nilpotent part of the Borel opposite to $\mathfrak{b}$, e.g. spanned by negative root spaces. Under the equivalence $\mathfrak{g}^* = \mathfrak{g}$, furnished by the Killing form, this identifies with $\mathfrak{n}$, the nilpotent part of $\mathfrak{b}$ itself, so we find that $T^*_{\mathfrak{b}}X = \mathfrak{n}$. Consequently T^*X is just another name for $\widetilde{N}$.

This switch in perspective opens some interesting doors (one could even say floodgates). That is because tangent bundles on manifolds carry extra structure: they admit a canonical symplectic structure. I don't really know how much you know about or care too deeply for symplectic geometry though.

4. Or a Lover of Classical Mechanics

The frame of mind I can adopt to care for symplectic manifolds is recalling how much fun classical mechanics is! (In such stark contrast to quantum mechanics[2].)

You've probably heard this spiel before, but here goes: what does one need to do physics?

i) Well, we're analyzing some phenomena, so there should be a phase space, encoded by a smooth manifold M, parametrizing all possible states that the mechanical system under consideration could be in. The functions on this phase space are physical quantities, which associate to every possible state of our system a value. We call these functions 'observables', betraying a belief that these are the things

[2]This opinion is from my December 2018 self who initially wrote this, and has since been thoroughly revised.

that we can, given a system in whatever state, observe, which is to say, measure.

ii) Because we are there and will measure something, we single out an observable $H : M \to \mathbf{R}$ and call it the Hamiltonian. It's supposed to correspond in a certain way to the 'total energy' of the system, but that's not too helpful. One POV is that it encodes our physical presence along the system (a POV that starts being important in QM). Another, perhaps saner POV, is that H is encoding all the laws of physics relevant to the experiment; that is certainly the approach one adopts when computing examples of mechanical systems. In any case, the Hamiltonian is neccessary, because before without it, we only have the phase space, the 'space of possible states of a system', without anything singling out the 'system' itself.

iii) So we have H, and it should tell us about the system itself that we are looking at. In particular, it should dictate the time evolution of the system; supposing it's in one state right now, where will be after a little time passes. You could object that perhaps that can't be determined, but the premise is that it can, i.e. that the universe is deterministic. This isn't that bizzare; a large goal of physics is making predictions, and actions eliciting reactions, etc., so such an assumption, at least on a very basic local level governing the behavior of mechanical systems, is always present in physics (from this point of view of what determinacy means, so far as I can tell even in quantum math).

Okay, that was a philosophical digression, justifying the claim that there should be a way of extracting time evolution out of the Hamiltonian. Time evolution should tell us for any state how the system will change from that state, and what state it has been in just prior; we are describing the notion of a flow, i.e. in classical terminology a '1-parameter family of deformations'. These come from a well-known source: from vector fields! Indeed, a vector field X on M defines locally (in coordinate patches) on M a first order PDE, which we can integrate given the initial data of a point in M, obtaining a flow on M. Conversely, the classical terminology calls a vector field the 'infinitesimal generator of the flow'.

So if we had a vector field at our disposal, we could integrate it to obtain a flow on M, which is in turn what encodes the time evolution of the mechanical system we are studying. But what we have (given ourselves so far) is just a function H on M. From it, we can extract a 1-form dH on M, which is a good step forward, but sadly it has the wrong variance. So we find ourselves requiring a mechanism to convert differential 1-forms into vector fields, i.e. a vector bundle map $T^*M \to TM$. This is equivalent to a pairing

$$\omega : TM \otimes TM \to \mathbf{R}$$

(on the level of bundles, e.g. $\mathbf{R}$ is the constant line bundle on M). We would like $T^*M \to TM$ to be an iso, or equivalently for the pairing of vector fields ω to be non-degenerate. In some ways, this ω is really the mechanism which encodes the laws of physics, or perhaps the immutable physical truths of the universe, as it allows us to transform the prescription for how the system is allowed to behave, i.e. the Hamiltonian, into the actual way in which the system does behave, i.e. evolve with time.

Then, giving ourselves such an ω, we can turn the 1-form dH into a vecor field, which we will denote X_H. The flow, $\phi_t : M \to M$ for $t \in \mathbf{R}$, of this field encodes the time evolution. What does this mean concretely? Well, given an observable $f : M \to \mathbf{R}$ (though the codomain could be any other smooth manifold too, if you wanted), which we think of as this observable at time 0, we define $f_t = f \circ \phi_t$ for any t to be what this observable will become after the passage of t time. Since f_t

represents the 'same observable' as f, only at a 'different time' let us follow the common convention and leave the subscript t implicit.

We can now also check how much a given observable f changes with time, i.e. form its derivative. Basic differential calculus manipulations show that

$$\frac{df}{dt} = X_H(f).$$

This gives some insight into the physical relevance of the Hamiltonian vector field X_H: if we view vector fields as derivations on functions, then X_H is the operator $\frac{d}{dt}$ of differentiating with respect to time. ♡

iv) The final remark sounds silly, but it's sort of important: the laws of physics should not change with time. What that means is that H, the observable by which we have gauged the passage of time, should be invariant under this time evolution. Since the latter is encoded by the flow ϕ_t of the vector field X_H, we require that H is constant wrt t. By what we know about $\frac{d}{dt}$, this is equivalent to requiring that $X_H(H) = 0$. Chasing through the definitions, we get

$$0 = X_H(H) = dH(X_H) = \omega(X_H, X_H).$$

This looks like a property of H, but at the end of the day, H was somewhat arbitrary. Indeed, we just chose something we can measure and gauges everything according to it, but we could well have chosen something else. Said differently, different observers might see different Hamiltonians, despite the phases being the same. This sounds like some relativity voodo, but it's really a more elementary observation. Though we first specified H and only then discussed ω, the latter should be something attached to the phase space M itself, and work the same for any choice of Hamiltonian. As said before, ω is encoding a much more fundamental physical nature than the choice of particular physical phenomenon that H specifies.

The takeaway is that we should require the equation $\omega(X_H, X_H) = 0$ to hold for all possible choices of Hamiltonians (= for all observables) H on M.

Though that isn't quite without loss of generality (unless M is simply connected, I think), let's instead assume we have a bit more: let's say that $\omega(X, X) = 0$ holds for all vector fields X on M. That is to say, ω is antisymmetric (as we're not in char 2), and as such can be identified with a differential 2-form on M.

v) The fundamental laws of nature (i.e. that structure of reality, which, unlike the laws of physics affecting only a particular system that were the subject of iii.), is independent of the particular physical system we are studying), does not change with time. Since those were incarnated in the 2-form ω on the phase space M, this requirement can be formalized as saying that $\frac{d\omega}{dt} = 0$. This should, just as M and ω are, be independent of the choice of Hamiltonian H inducing the time evolution. That is to say, it should hold for all choices of $H : M \to \mathbf{R}$.

Recall that the operator $\frac{d}{dt}$ of time derivative in the physical system determined by a Hamiltonian $H : M \to \mathbf{R}$ is given by the Hamiltonian vector field X_H. According to a basic but useful result of differential calculus on manifolds, literally called Cartan's magic formula, we have

$$\tfrac{d}{dt}(\omega) = (d\omega)(X_H, -) + d(\omega(X_H, -)).$$

By definition of the vector field X_H, we have $dH = \omega(X_H, -)$, so the second term on the rhs above vanishes by $d^2 = 0$. Time invariance of ω is thus equivalent to vanishing of $d\omega$, evaluated at Hamiltonian vector fields on one variable. Since the choice of H was arbitrary above, this is equivalent to $d\omega$ vanishing when evalated by any vector field on one variable, and hence to $d\omega = 0$. So ω must be a closed 2-form.

And we're done! We've pursued the bare minimum setup needed for physics, and we have ended up with a smooth manifold M with a closed non-degenerate 2-form ω - a symplectic manifold!

The takeaway is the following: symplectic geometry may be viewed as the study of Hamiltonian classical mechanics. Such a perspective of course isn't always beneficial, but sometimes (or at least for me, basically always) it sheds some light on the various constructions one encounters in symplectic geometry.

Anyway, for what proports to be physics, the last few paragraphs were rather abstract, so let us see how the classical mechanics that we all know and love (in the Hamiltonian setup) constitutes an example of the above paradigm.

Give a mechanical system, it will possess a space of possible positions, the configuration space in the form of a smooth manifold M. This isn't the phase space, as it's just the configurations of the system; a dampened pendulum might return to the same position a number of times, but each time with a different velocity, so it attain the same configuration at different phases. So what else do we need to specify to know the state of the system, other than its position in space? Newton's law can be understood as saying that the position and velocity determine all other physical attributes of a system, hence we need to add in velocities. If q is a coordinate on M, then surely its velocity would most naturally be the tangent vector $\frac{d}{dq}$, but it turns out to be more convenient to encode it as the 1-form dq instead. That is to say, we take the cotangent bundle T^*M to be the phase space. The Hamiltonian will be any smooth function $H : T^*M \to \mathbf{R}$. Classically you would take $H = T + U$ for T the kinetic energy and U some potential.

There is one more piece of data needs: the symplectic form ω. In order to find it we work locally, letting q be a coordinate on M, which you may view as 'horizontal' coordinates on T^*M, and $p = dq$ be the corresponding fiber or 'vertical' coordinates of T^*M. Together p and q form a coordinate system on T^*M, in which we define the 2-form $\omega = dp \wedge dq$ (note that because $p = dq$ isn't really itself a differential, only a coordinate on T^*M, its differential dp needn't be zero and all is fine). The claim is that this always defines a symplectic form on the phase space T^*M.

So woohoo, we've equipped an arbitrary cotangent bundle with a symplectic structure (as promised in the previous section), but we've done so 'accidentally' or 'along the way' of motivating ourselves about symplectic manifolds in the first place!

In fact, the cotangent bundle is a rather generic example of a symplectic structure; a theorem of Darboux guarantees that any symplectic manifold is localy isomorphic to a cotangent space with its canonical symplectic structure. The local coordinates realizing such an iso are called called canonical coordinates, and also play a prominent role in classical Hamiltonian mechanics. So the fact that we usually do mechanics with a specified configuration space inside the phase space is not a severe restriction - working locally in canonical coordinates that can always be arranged.

Before we move on, let's point out that the time evolution law on T^*M with a Hamiltonian H takes on an especially familiar form. In terms of the usual canonical coordinates (q, p) on T^*M described above, a calculation shows the associated Hamiltonian vector field to be

$$X_H = \frac{dH}{dp}\frac{d}{dq} - \frac{dH}{dq}\frac{d}{dp}.$$

Thus time evolution of the observables p and q (and since they are the coordinates on the phase space, it really suffices to specify just those) is hence expressed as

$$\begin{array}{ccccc} \frac{dq}{dt} & = & X_H(q) & = & \frac{dH}{dp}, \\ \frac{dp}{dt} & = & X_H(p) & = & -\frac{dH}{dq}, \end{array}$$

recovering Hamilton's equations, in a form in which you've surely encountered them in some physics class.

5. Hamilton was a Man of Action

Before we return to the flag variety and ask what the symplectic techniques can do for us there (the answer, unsurprisingly, is quite a lot), let's spend a little more time on symplectic geometry = classical mechanics.

In physics, there is often a lot of talk of symmetry. A symmetry of a mechanical system is a group action which preserves everyhing in such a way, that the physics remains unaffected. Finding symmetries of a system is a crucial technique for working out essentially any more involved examples, the spinning top being perhaps the most famous. The primary thing symmetries buy you is constants of motion (that's the celebrated Noether's Theorem, the one great thing Our Lady did outside abstract algebra), which are terrifically useful to have.

Technically speaking, from a Hamiltonian perspective on mechanics that we are pursuing, a symmetry of a system (M, ω, H) is a smooth action of a Lie algebra G on M, which satisfies some compatibility with the symplectic structure, and for which the Hamiltonian function H is G-invariant. The compatibility condition, which goes by the name Hamiltonian action in symplectic geometry, comes in two parts:

- Firstly, G must act on M through symplectomorphisms, which is to say that for any $s \in G$, the translation map $s : M \to M$ preserves the symplectic structure ω.

- Secondly, the associated infinitesimal action, which means the induced Lie algebra map $a : \mathfrak{g} \to \{\text{vector fields on } M\}$ is required to land inside Hamiltonian vector fields. Said differently, for every $x \in \mathfrak{g}$, the vector field $a(x)$ on M must be of the form $d(H_x)$ for some function[3] H_x on M.

Summarizing informally, the group G acts on the phase space is such a way, that the infinitesimal generators of this action are the 'physically meaningful' ones, i.e. the ones which give rise to time evolution wrt some Hamiltonian. Equivalently, the curve $t \mapsto \exp(tx)m$ for a fixed $x \in \mathfrak{g}$, fixed $m \in M$, must be a trajectory of the state m under the motion dictated by some Hamiltonian. From the POV of mechanics, that is a very reasonable thing to assume!

Btw, the constants of motion promised by Noether's Theorem are precisely the observables H_x on M for $x \in \mathfrak{g}$. Said theorem is therefore quite straightforward from this

[3]The reqirement that $a(x) = d(H_x)$ specifies the observable H_x only up to a constant (or pedantically, a locally constant function). Picking the constants for H_x simultaneously for all x in an appropriately uniform way, making the induced moment map $\mu : \mathrm{Obs} \to \mathfrak{g}^*$ (we'll talk about it much more in what follows) continuos isn't all that difficult.

But it also doesn't specify the map μ (or equivalently, all those constants) uniquely yet: there is still one constant of room. To determine it (i.e. set some kind of initial value), recall the requirement that the moment map be G-equivariant. Under at least some assumptions (I forget the precise details, but I think either if G is compact or semisimple, we're good - otherwise there is a correction of equivariance by a $\mathfrak{g}^*$-valued cocycle game that one needs to play), this serves to uniquely pin down the constants, and with that, allows us to talk about "the" moment map.

post-modern POV on mechanics - it amounts to the claim that $\frac{d(H_x)}{dt} = 0$, which is a simple calculation.

Given a Hamiltonian action of G on a symplectic manifold (M,ω), we obtain the so-called moment map $\mu : M \to \mathfrak{g}^*$ by sending a point m to the functional $x \mapsto H_x(m)$ on the Lie algebra $\mathfrak{g}$. (Actually, the name moment map is an unfortunate mistranslation from French, and it should be called the momentum map for physics reasons. Alas, the name seems to have stuck.)

Why the name moment(um) map? Well, in the arguably most classical mechanical setup of a point in 3-space, with phase space $M = T^*(\mathbf{R}^3)$, the rotarional symmetry of the underlying phase space $\mathbf{R}^3$ induces a Hamiltonian SO(3)-action on the phase space. The associated momentum map $\mu : M \to \mathfrak{so}(3) = (\mathbf{R}^3, \times)$, where $\times$ denotes the cross product, identifies precisely with the angular momentum from basic classical mechanics: in standard canonical coordinates $\mu(q,p) = q \times p$. If state (q,p) is realized by a mechanical system, q is the position vector and $p = \frac{dq}{dt}$ is the velocity, so it's really the angular momentum you remember!

The moment map is an oft-used powerful tool in symplectic geometry; for instance, under some mild extra conditions, the preimage $\mu^{-1}(0)$ inherits a G-action from M, and its quotient by this action is what's called the symplectic quotient of M by G, or the Hamiltonian reduction of M. The ever-so slightly unfortunate, but completely standard, notation for it is $M//G$.

Of course there is nothing special about the point $0 \in \mathfrak{g}^*$, and we could analogously form variants of Hamiltonian reduction $\mu^{-1}(x)/G$ for any regular value x of the moment map μ.

If you wish to avoid the pesky niceness requirements on the G-action, you know what to do: go derived, form homotopy quotients instead of ordinary ones, rake in the profits! Pavel Safranov has written some exciting papers on derived Hamiltonian reduction. In fact, he does it even in the shifted symplectic setting, so it's even more exiting!

6. I am the Moment, the Moment is Me

We've gotten so far, wading neck-deep through the dense swamps of symplectic geometry = classical mechanics. Wasn't this supposed to be all about rep theory?! Can't we just stop, turn back?

No!! We've come this far, no sense giving up now! I am not throwing away my shot!! At explaining how all this ties back to the Springer resolution.

In hopes of understanding Hamiltonian actions better, let's go back to the one sort of symplectic manifolds we understand, the contangent bundle T^*M, i.e. back to the classical mechanics where we have specified configurations inside phases. Then a particular case of a Hamiltonian G-action on T^*M is any action induced on the cotangent bundle by an action of a Lie group G on the underlying smooth manifold M. So any symmetry of the configuration space instantly produces a symmetry of the whole mechanical system.

In this case, Hamiltonian reduction turns out to be $T^*(M/G)$, at least when G acts nicely (free and prop discont, say) on M. In symbols $(T^*M)//G = T^*(M/G)$, so the symplectic quotient "goes inside T^*". The moment map $\mu : T^*M \to \mathfrak{g}^*$ also admits an explicit description: its associated map of functions (let's be in AG again, just for a moment) $\mathrm{Sym}^*(\mathfrak{g}) = \mathcal{O}(\mathfrak{g}) \to \mathcal{O}(T^*M) = \Gamma(M; \mathrm{Sym}*(T_M))$ is generated as an algebra map by the linear map

$$\mathfrak{g} \xrightarrow{a} \Gamma(M; T_M) \to \Gamma(M; \mathrm{Sym}^*(T_M)),$$

with a the infinitesimal action and the second map induced by the inclusion $T_M \to \mathrm{Sym}^*(T_M)$.

Most importantly for our purposes, this means that the left G-action on G/B (or equivalently, the conjugation action on the Borels) gives rise to a Hamiltonian G-action on the cotangent bundle of our good-old flag variety; a Hamiltonian G-action on $T^*(G/B) = \widetilde{N}$.

Wherever there is a Hamiltonian action, there is a moment map, which in this case maps $\widetilde{N} \to \mathfrak{g}^*$. Post-composing with the Killing form isomorphism, we get a map $\widetilde{N} \to \mathfrak{g}$, which is precisely the Grothendieck resolution map, i.e. the resolution of singularities. People who symplect to survive know a lot about moment maps, so the mere fact that this is a moment map allows you to conclude various things about the Grothendieck resolution.

In §3, I implied that discovering the Grothendieck resolution in the guise of $T^*(G/B)$ could happen accidentally. Hopefully §3-6 illustrated exhaustively that conversely one could also come upon this incarnation of $\widetilde{N}$ knowing what they were doing very well indeed.

6.3/4. A Brief Word From Our Sponsors

Before departing from the symplectic realm, now that we've seen what we came here to see, allow me to just mention one other important rep theoretically significant occurence of symplectic stuff.

The relevant thing is the Riemann-Hilbert Correspondence. I know you have been told to not care too deeply about it, hence I'm relegating it to a passing mention. RH is at its essence a vast generalization of the classical differential geometric equivalence between vector bundles with flat connections and local systems (= locally constant sheaves of vector spaces). It gives an equivalence between perverse sheaves and regular holonomic D-modules. Perverse sheaves are a very useful generalization of constructible sheaves, about which I want to say nothing other than that they, along with the RH correspondence, were used to solve (or better, to reduce it to things trivial in light of the Beilinson-Bernstein thm) the Kazhdan-Lusztig Conjecture. What I do wish to point out, on the other hand, is that the condition of holonomicity on D-modules is something that has to do with symplectic manifolds.

Very briefly: for a coherent D-module $\mathscr{F}$ on a smooth variety X, there exist good filtrations, compatible with the degree filtration on the sheaf of differential operators D_X, and s.t. the associated graded modules are each finitely generated (sadly there isn't a canonical choice for such a good filtration, but luckly everything ends up being independent of its choice anyway). Then $\mathrm{gr}(\mathscr{F})$ is a coherent sheaf of $\mathrm{gr}(D_X) = \mathrm{Sym}^*(T_X)$-modules, or equivalently a coherent sheaf on the cotangent bundle T^*X. Taking its support produces a closed subvariety of T^*X, denoted $\mathrm{SS}(\mathscr{F})$ and called either the characteristic variety (in analogy with characteristics in differential equations) or the singular support of $\mathscr{F}$.

This $\mathrm{SS}(\mathscr{F})$ isn't just any old subvariety of T^*X. The canonical symplectic form omega vanishes on it, making singular suppored a so-called isochronic submanifold. Sadly I don't have a very good understanding of what this condition means on the level of mechanics. But one thing I can tell you about them that their dimension is bounded below by one half $\dim X$ (recall that any symplectic manifold has to be even dimensional, since odd dimensional vector spaces can't support non-degenerate antisymmetric bilinear pairings by easy linear algebra). If this lower bound is achieved, the isotropic submanifold is called a Lagrangian submanifold.

Now Lagrangian submanifolds, they are the stars of symplectic geometry. Their intersection theory is superbly well-behaves and who knows what else, but most crucially (for me to care), they admit an interpretation in terms of mehcanics. Indeed, given the symplectic manifold T^*M, the phase space of a classical mechanical system, the zero section M in T^*M is an example of an isotropic submanifold. Since $\dim T^*M = 2\dim M$, the zero section is actually Lagrangian. And in canonical coordinates, any Lagrangian submanifold is of this form. That is to say: a Lagrangian submanifold in a symplectic one corresponds to specifying a configuration space inside the phase space.

A coherent D-module is by definition holonomic iff its singular support is Lagrangian. So you see a situation where the symplectic geometry plays a crucial role in a purely D-module story. There's more to say about this, but let's not. We've dwelt in symplecticland long enough!

7. Why are We Neglecting Poor Alexander

At this point, symplectic and mechanics mussings will probably take up more than half of this text. But it's really all about rep theory, namely about the Springer resolution $\widetilde{N}$!

Well, maybe not all; the title mentions Grothendieck too. So let's focus on his contribution to the story.

Seeing the great success we had with the Springer resolution by defining it as an incidence set of borels and their elements, we could try an even more basic version of the same idea:

$$\widetilde{\mathfrak{g}} = \{(x, \mathfrak{b}); \mathfrak{b} \text{ a Borel in } \mathfrak{g},\ x \text{ a point in } \mathfrak{b}\}.$$

This is the Grothendieck-Springer resolution, sometimes called the GS simultaneous resolution or the Grothendieck alteration. It has an evident map to the flag variety $X = G/B$ by projecting onto the second factor, and an almost obvious one to the universal Cartan $\mathfrak{h}$, comprised of projecting onto the first factor, composed with the universal quotient map $\mathfrak{b} \to \mathfrak{h}$.

In a very similar spirit, there is a map $\widetilde{\mathfrak{g}} \to \mathfrak{g}$ given by $(x, \mathfrak{b}) \to x \in \mathfrak{b} \subset \mathfrak{g}$. This is the analogue of the Springer resolution map $\widetilde{N} \to N$, but also its generalization, since $\widetilde{N}$ is contained in $\widetilde{\mathfrak{g}}$, and the restriction of the GS resolution map is precisely the Springer resolution map.

The last two maps from the GS resolution are related by commutative square

$$\begin{tikzcd} \widetilde{\mathfrak{g}} \arrow[r] \arrow[d] & \mathfrak{h} \arrow[d] \\ \mathfrak{g} \arrow[r] & \mathfrak{h}//W \end{tikzcd}$$

Here the right vertical arrow is the quotient projection, the lower horizontal arrow is the Chevalley map, and the other two arrows have just been discussed above. The commutativity of the diagram should basically be obvious, in light of how the Chevalley can either be expressed in terms of restriction to a chosen embedded Cartan, or in terms of the quotient map to the universal Cartan.

While commutative, the above square only becomes a pullback if we restrict to regular elements of $\mathfrak{g}$ everywhere. Indeed, if I'm understanding everything right, an element of $\mathfrak{g}$ is regular iff it's a regular point (in the analysis sense) for the Chevalley map.

That the square is Cartesian over the regular locus is rather easy to see though (disclaimer: this is still something of a proof-sketch, as this is only an note. But basically everything is there). The claim is that $\widetilde{\mathfrak{g}}_{\text{reg}} \to \mathfrak{g}_{\text{reg}}$ is a W-bundle[4], and that the map $\widetilde{\mathfrak{g}}_{\text{reg}} \to \mathfrak{h}$ is W-equivariant. In that case the square encodes a map of W-bundles, and since the fibers are the same (what with both being W on account of being W-bundles[5]), the maps between them are isos (which is itself equivalent to Cartesianness of the square). But a fiber over a regular x is Borels containing the Cartan $\mathfrak{h} = C_{\mathfrak{g}}(x)$, i.e. the choices of which roots are the positive roots in a fixed root lattice. Since the Weyl group is, by its definition in terms of the root data, precisely the full symmetry group that permutes these choices among each other, and clearly does so simply transitively, the claim follows.

The above square also contains a lot about $\widetilde{N}$ and N. Indeed, N may be seen as the fiber of the lower horizontal map $\mathfrak{g} \to \mathfrak{h}//W$ at the trivial coset, by a description of N that we discussed near the end of §1 together with the Chevalley Restriction Theorem. Similarly, but more tautologically, $\widetilde{N}$ is the fiber of the upper horizontal map $\widetilde{\mathfrak{g}} \to \mathfrak{h}$ at 0. Therefore the pullack of this square along the inclusion of 0 on the right produces the Springer resolution map $\widetilde{N} \to N$. In particular, knowing that the square is Cartesian over the regular locus and that the fibers W are finite means that we have proved in a way that is substantially less sketchy way than before that the Springer map is actually birational (there's also the matter of the codimension of $\mathfrak{g}_{\text{reg}}$ in $\mathfrak{g}$ that I'm sweeping under the rug here - actually, a very nice proof of that also relies on simplicial structures on coadjoint orbits and moment maps! See the end of §10 for a little more on that.).

Btw, note how elegantly $\widetilde{\mathfrak{g}}$ allows us to express the Weyl group, as the fiber of $\widetilde{\mathfrak{g}}_{\text{reg}} \to \mathfrak{g}_{\text{reg}}$. Contrast with the usual way one defines W, either by reflections in the weight lattice, or by picking a maximal torus H in a connected group G with Lie algebra $\mathfrak{g}$ and forming $W = N_G(H)/H$. Both require making choices, of $\mathfrak{h}$ and $\mathfrak{b}$ or of G and H respectively. Technically the GS approach also requires making a choice of the regular element at which to take the fiber, the Weyl group has an independent role as the structure group of the torsor. Anyway, apparently this POV is useful for certain things (which I don't understand) some people like to do with Weyl groups.

8. Maximizing Individual (Instead of Social) Utility

So why do we, perhaps (for the matter of argument) fancying ourselves geometric representation theorists, care? Well, let's recall where the Grothendieck-Speinger resolution enters the proof of Beilinson-Drinfeld Localization. (Surely you know all this at least as well as me if not much better, but allow me to include it just for completeness.)

The setup is such: you have already proved that the structure map $U(\mathfrak{g}) \to D(X)$ for $X = G/B$ the flag variety factors through the quotient map $U(\mathfrak{g}) \to U(\mathfrak{g})_0$, moding out the kernel of the central character. We wish to show that the induced map $U(\mathfrak{g})_0 \to D(X)$ is an iso.

To do this, you do what else but pass to associated graded wrt the PBW and order filtrations respectively. We compose this map of associated gradeds on both sides with respectively the surjection

$$\mathcal{O}(N) = \operatorname{Sym}^*(\mathfrak{g})/\operatorname{Sym}^+(\mathfrak{g})^G \to \operatorname{gr}^*(U(\mathfrak{g})_0),$$

[4]I was alerted to the fact that there is some subtlety involved in what we mean by a "W-cover" here. In particular, *it is not* correct to say that the map in question is a W-torsor. But let us turn a blind eye, and pretend all is fine.

[5]There is some more of the same sort of cheating going on here. Actually, what we said only holds over the so-called regular semisimple elements of $\mathfrak{g}$, over which we do genuinelly get W-torsors. Extending that to the whole of $\mathfrak{g}_{\text{reg}}$ take some work.

coming from the PBW theorem and the identification $\mathrm{gr}^*(Z(\mathfrak{g})) = \mathrm{Sym}^*(\mathfrak{g})^G$, and the injection

$$\mathrm{gr}^*(D(X)) \to \mathrm{Sym}^*(T_X) = \mathcal{O}(\widetilde{N})$$

coming from the "PBW for D_X" and left exactness of taking global sections. The isos with functions are things we already encountered along the way, in §1 and §3 respectively. Chasing through the definitions shows the map of functions to be the one indiced by the Springer resolution map $\widetilde{N} \to N$. If we know Kostant's result that it's birational (and an iota more), it follows that it induces an iso $\mathcal{O}(N) \to \mathcal{O}(\widetilde{N})$. But then the composite

$$\mathcal{O}(N) \xrightarrow{\text{sur}} \mathrm{gr}^*(U(\mathfrak{g})_0) \to \mathrm{gr}^*(D(X)) \xrightarrow{\text{inj}} \mathcal{O}(\widetilde{N})$$

is a bijection, so the leftmost map must in particular be injective and the rightmost map likewise surjective. Both are therefore bijective, proving the middle map to be a bijection also. But the functor gr^* is conservative, so indeed $U(\mathfrak{g})_0 \to D(X)$ is an iso.

9. You Can Still be a Differential Geometer

So far we've seen the utility of the GS resolution, but only one construction of it. Recall that, other than as an incidence set, the Springer resolution can also be realized as $T^*(G/B)$. What's the analogous cotangent bundle-level description of $\widetilde{\mathfrak{g}}$?

Let's start with $Y = G/N$, the 'extended flag variety' or the 'base affine' (a piece of terminology I find very perplexing). The inclusion $N \subset B$ (as its unipotent radical, no less) induces a map $G/N \to G/B$. The fiber is $(G/N)/(G/B) = B/N = H$, so we have exhibited $Y \to X$ as an H-torsor. In particular, Y carries a narural H-action wrt which $Y/H = X$. Note though that this is not the H-action restricted from the obvious left G-action on G/N, at least there is no reason for the two actions to agree so far as I can tell.

The space Y with its H-action is quite important rep theoretically. For instance, twisted D-modules on X can be profitably viewed as weakly H-equivariant D-modules on Y, together with specified monodromy, to be given by the character we're twisting by, along the H-orbits. (This is only about 89% clear to me right now, but I'm getting there. If you know how to rigorously prove this, I'd be very happy if you could show me sometime.) Supposing the twisting was trivial, this would for instance produce (strongly) H-equivariant D-modules on Y, which is of course the same thing as D-modules on X, which is certainly what non-twisted D-modules shoud be.

Relatedly, if you wish to have a Beilinson-Bernstein Localization Theorem realize $\mathfrak{g}$-modules as D-modules on something like the flag variety, but you don't wish to specify the central character, one thing you could do is consider all weakly H-equivariant D-modules on Y. This is discussed in one of Ben-Zvi and Nadler's papers, and it turns out that this sort of localization is an equivalence, modulo some jazzed up Weyl group action.

Enough advertizing for Y, let's see how it connects to the Grothendieck-Springer resolution. Consider the cotangent bundle $T^*Y \to Y$. Its fiber over a point $y \in Y$, i.e. the cotangent space to Y at y, is easily seen to be $\mathfrak{g}^*/\mathfrak{n}^* = \mathfrak{b}$ by standard Killing form identification of duals, where $\mathfrak{b}$ is the Borel in $\mathfrak{g}$ that is the image of y under the projection map $Y \to X$. Thus T^*Y is seen to consist of pairs (y, x) of a point $y \in Y$ (whatever that really is) together with an element x of the Borel subalgebra $\mathfrak{b} \subset \mathfrak{g}$ determined by y.

Note that the H-action on T^*Y, naturally extending the H-action on Y, is trivial along fibers. Indeed, it would act by $\mathrm{ad}(s)$ in the fiber direction for $s \in H$, but conjugation in abelian groups isn't very interesting. This means that the quoient $(T^*Y)/H$ becomes identified with pairs $([y], x)$ for $[y] \in Y/H = X$ and x an element of the corresponding Borel $\mathfrak{b}$ - and we've ended up finding a cotangent bundle description of $\widetilde{\mathfrak{g}}$!

I think comparing the formulas $\widetilde{N} = T^*X$ and $\widetilde{\mathfrak{g}} = (T^*Y)/H$ really highlights both the comparison and the differences between the two constructions.

10. Or a Fan of Mechanics, Of Course

Writing $\widetilde{N}$ and $\widetilde{\mathfrak{g}}$ like that, our spider sense start tingling. Recalling that T^*Y carries a symplectic structure, and that the H-action on Y equips T^*Y with a Hamiltonian H-action, we are reminded of a formula for Hamiltonian reduction that we saw in §6.

Indeed, the (wlog Killing iso composed) moment map $\mu : T^*Y \to \mathfrak{h}^*$ maps, in terms of the discussed incidence description, explicitlyy u,

$$(y, x) \to y \in \mathfrak{b} \to \mathfrak{b}/\mathfrak{n} = \mathfrak{h}.$$

It clearly folows that $\mu^{-1}(0)/H$ is nothing but the incidence description of $\widetilde{N}$. In fancy words, we have realized the Springer resolution as the Hamiltonian reduction of the Grothendieck-Springer resolution. In fancy symbols: $\widetilde{N} = \widetilde{\mathfrak{g}}//H$.

In fact, we could have taken any other point $x \in \mathfrak{h}$ too (quite possibly I should also assume that x is regular throughout), rather than the origin 0 as we did above. The variant of the Hamiltonian reduction obtained that way (the action of H on T^*Y is good enough that there is no need to worry about anything) forms a resolution of singularities for the associated fiber of the map $\mathfrak{g} \to \mathfrak{h}//W$ at the corresponding coset (this follows from the exact same considerations as for $x = 0$, namely restricting to regular elements and pulling back the square from §7).

We can be more explicit about what this latter fiber over the oribit of x is. In this setting, it is more common to resist the urge to post-compose the moment map with the Killing iso as we had some before, so let us observe this convention. In particular, we are looking at the fiber of the map $\mathfrak{g}^* \to \mathfrak{h}^*//W$ at the orbit of an element $x \in \mathfrak{h}^*$. By Chevalley, it is the same thing as the orbit of the quotient map $\mathfrak{g}^* \to \mathfrak{g}^*//G$ at the orbit determined by x, i.e. the coadjoint orbit that contains x. Again by Chevalley, any coadjoint orbit contains some such x.

We may conclude that restricting the (Killing iso composed) GS resolution $\widetilde{\mathfrak{g}} \to \mathfrak{g}^*$ to any coadjoint orbit $\mathcal{O}$ in $\mathfrak{g}^*$ contains a resolution of singularities. This is the sense in which $\widetilde{\mathfrak{g}}$ is a 'simultaneous' resolution. Furthermore the subvariety of $\widetilde{\mathfrak{g}}$ which is the resolution of a given coadjoint orbit $\mathcal{O}$ in $\widetilde{\mathfrak{g}}$ can be determined quite concretely: picking an element $x \in \mathcal{O} \cap \mathfrak{h}^*$, it consists of all the points in $\widetilde{\mathfrak{g}}$ of the form $(\mathfrak{b}, x + y)$ for $x + y \in \mathfrak{b}$ and $y \in \mathfrak{g}$ nilpotent.

The reason we rather didn't Killing iso the difference between $\mathfrak{g}$ and $\mathfrak{g}^*$ away in this section, is similiar to why we prefer to phrase Hamiltonian mechanics in terms of the cotangent bundle as opposed to the tangent bundle. Indeed, the coadjoint orbits come equipped with a symplectic structure (even more: the obvious action of G is Hamiltonian wrt this symplectic structure) as an byproduct of the above discussion. But in fact, this symplectic structure, called the Kirilov-Kostant symplectic structure, is quite classical and admits a nice naive definition:

The G-action on a fixed coadjoint orbit $\mathcal{O}$ being transitive implies that the infinitesimal action map $a : \mathfrak{g}t\o\{\text{vector fields on } \mathcal{O}\}$ (which is in this case nothing but the lowercase coadjoint action ad^*) is too. For any x belonging to $\mathcal{O}$, and any $u, v \in \mathfrak{g}$, it thus suffices to define

$$\omega_x(a(u), a(v)) = x([u, v]),$$

where on the rhs we have evaluated x as the linear functional on $\mathfrak{h}$ that it is. The claim is that the 2-form ω on $\mathcal{O}$ is a symplectic form.

The very exietence of a symplectic structure on coadjoint orbits is very magical. Since only even dimensional manfiolds admit symplectic structures, it follows that all coadjoint orbits are even dimensional. From this, it for instance instantly follows that N_{reg} in N has codimension one, since $N - N_{\text{reg}}$ is a union of adjoint orbits with at least one missing, and as such has codimension at least 2. (Okay, I guess you still needed to have already known that N contains finitely many orbits..).

11. Maybe You're a Fan of Borel-Bott-Weil

As the Monthy Pythons were fond of saying: and now for something completely different! And by different, just as the Monthy Python, I of course mean a different way of approaching the Grothendieck-Springer resolution.

Recall the Borel-Bott-Weil Theorem, on irreducible highest weight representations as G-equivariant vector bundles on the flag variety $X = G/B$. If the Beilinson-Bernstein Localization Theorem marks the beginning of the field of geometriy representation theory, as I have often heard people assert, then surely Borel-Bott-Weil is the immediate prehistory, which layed the groundwork, fermented wild new ideas in public imagination, and made the air brimming with anxiety and anticipation for the inevitable revolution.

Anyhow, to a B-rep M, you associate the vector bundle $L(M) = G \times_B M$, where the subscript G stands not for fibered, but for balanced product (this has some potential to be confusing. I have seen some people denote balanced products by superscripts, which might be the way to go). The terminal projection $M \to \text{pt}$ gives rise to the bundle projection $L(M) = G \times_B M - - > G \times_B \text{pt} = G/B$, and the fiber is clearly M. The content of BBW is that this allows you to convert representations into vector bundles and study those, to great avail as you now have the full power of algebro geometric theory of sheaves and bundles at your disposal. Especially highest weight reps become especially accessible, as we know a lot about line bundles on varieties.

So let's do the BBW for $M = \mathfrak{b}$, with the adjoint action. We get $L(\mathfrak{b}) = G \times_B \mathfrak{b}$, yielding a vector bundle with fiber $\mathfrak{b}$ on X. Let's think about what we're doing here: we take pairs (s, x) of an element $s \in G$ and $x \in \mathfrak{b}$, and identify for every $t \in B$ the pair (st, x) with $(s, \mathrm{Ad}(t)x)$. In that way the first factor traverses the elements of the poset $sB \in G/B$. Fixing the value of sB (but not of s itself), the second factor traverses the Borel subalgebra $\mathrm{Ad}(s)\mathfrak{b}$ of $\mathfrak{g}$ (since the adjoint action of B preserves its Lie algebra $\mathfrak{b}$, this is independent of the choice of coset representative s). This is the Lie algebra of the Borel subgroup sBs^{-1} in G. Thus we recognize once again the incidence description of $\widetilde{\mathfrak{g}}$.

This description of the GS resolution is in some ways inferior to the previously discussed ones, since it requires choosing a Borel subgroup $B \subset G$. Note that before, even though we often denoted the flag variety as G/B, we did not actually need to choose B, since the flag variety can well be identified with the space of Borels.

One genuinely interesting thing to be gained by viewing $\widetilde{\mathfrak{g}}$ this way though is that it makes evident the equivalence of quotient stacks $\widetilde{\mathfrak{g}}/G = \mathfrak{b}/B$, where the rhs is quotient wrt to the adjoint action. This might seem minor, but phrasing rep theory in terms of algebraic geometry, these sort of objects come up a lot and encode much.

For instance, much of the highest weight yoga is encoded in the correspondence of adjoint action quotient stacks

$$\mathfrak{g}/G \leftarrow \mathfrak{b}/B \rightarrow \mathfrak{h}//H,$$

induced by choosing a Borel subgroup $B \subset G$. But by the above observation, we may replace the midele term with $\widetilde{\mathfrak{g}}/G$, removing the seeming dependence on any choices. I think this, its appearence in the above fundamental correspondence, probably also best highlights the importance of $\widetilde{\mathfrak{g}}$ for rep theory.

12. Quantization is a Fancy Word and Little More

To conclude this massive exposition dump on the Springer and Grothendieck-Springer resolutions, I would like to paraphrase and elaborate upon an answer that David offered to a mathoverflow question about the importance of the Springer resolution.

It has to do with quantization. Now I don't know how you are, but I for one freeze wherever anyone mentions that word. But for our purposes here, quantization won't be anything scary. It will just mean the following: we'll be given a graded commutative algebra A. Its quantization will mean just a usually non-commutative filteted algebra $\hat{A}$ for which $\mathrm{gr}^*(\hat{A}) = A$. We'd also say that A is the classical limit of $\hat{A}$. We might also say that $\hat{A}$ is the quantization of $\operatorname{Spec} A$, following the paradigm that the classical role of spaces is played in the quantum realm by "noncommutative spaces".

The prime example is $U(\mathfrak{g})$, which we view as a quantization of $\mathrm{Sym}^*(\mathfrak{g})$, or by Killing iso of $\mathfrak{g}$ itself. Similarly the sheaf of differential operators D_X is viewed as a quantization of the cotangent bundle T^*X (and if the latter is where classical mechanics happens, then ... wink wink, nudge nudge!).

What does the Springer resolution, or more precisely its factorization

$$T^*X = \widetilde{N} \to N \to \mathfrak{g},$$

quantize to? Well, we've already seen that $\mathcal{O}(N)$ is the associated graded wrt the PBW filtration of the quotient $U(\mathfrak{g})$ of the enveloping algebra by the kernel of the trivial character. Hence the above springer factorization quantizes to

$$U(\mathfrak{g}) \to U(\mathfrak{g})_0 \to D(X).$$

The fact that the Springer resolution $\widetilde{N} \to N$ is birational, i.e. that it induces an iso on functions (and it is in some sense really the algebras of function that we are quantizing), reflects that the above factorization induces an iso $U(\mathfrak{g})_0 = D(X)$.

This is the sense in which the Springer resolution quantizes to the Beilinson-Bernstein theorem (though of course, the full theorem identifies more than just algebras, namely of categories). In fact, this (perhaps with fewer mentions of quantization) is the usual argument one uses in the course of proving BB, and certainly the one we discussed in §8.

A phrase one sometimes hears is that "Quantization is not a science, but an art!", which is to say that the same classical phenomenon may admit several different quantizations (and finding the most physically relevant one might not be easy). Despite our (you and me, right here and now), shall we say, heavily curtailed understanding of what quantization entails (- passing to the associated graded), we can see this very well: several different filtrations may produce the same associated graded.

Indeed, $\widetilde{N} = T^*X$ could, instead of as $D(X)$, also be quantized as $D^\lambda(X)$, the ring of λ-twisted differential operators, for any dominant (I was assured that this is the right word to say) weight $\lambda \in \mathfrak{h}^*$. In that case, to remain compatible, we should also quantize N not as $U(\mathfrak{g})_0$, but as $U(\mathfrak{g})_\lambda$. (Really that's nothing scary, just determining the associated gradeds of these filtered algebras.) With this choice of quantization, Springer resolution quantizes to the version of the BB theorem with the central character associated to lambda.

The takeaway can therefore be: "The Springer resolutoion quantizes to Beilinson-Bernstein Localization."

I could be mistaken, but I am willing to wager a guess that the Grothdieck-Springer resolution analogously quantizes to the "Beilinson-Bernstein over the base affine" version from David's paper, i.e. version for weakly H-equivariant D-modules on $Y = G/N$, mentioned in §9.

With that piece of unfounded bs speculation, it is high time to wrap up this write-up.