Our Mathematical Senses

The Geometry Vision

Prof. Vijay Ravikumar

Department of Mathematics

Indian Institute of Technology- Madras

Lecture-63

Video 12D: visualizing the action of PGL(3,R)

So let's visualize the action of PGL3R on RP2. So remember, RP2 is the set of lines through the origin in R3. And we want to know the action of PGL3R on those lines. PGL3R takes lines through the origin to lines through the origin. So it will act on RP2. But it's a little hard to visualize all that.

So we're going to use this affine chart. Remember, there's an affine chart we defined called phi z, which takes RP2 to the plane z equals 1. And the way it works is it'll take any line through the origin here and map it to this point in z equals 1 where it hits z equals 1. So we can use this affine chart, this plane z equals 1, to visualize most of RP2.

So using that, let's look at the action of PGL3R on RP2 via the induced action on this chart z equals 1. So here's an element of PGL3R. Here's an element of this plane z equals 1. You can see that the z-coordinate is equal to 1. Where does this get sent? By this matrix.

Well, we just do matrix multiplication. We multiply this vector by this matrix. That'll give us a new vector. That'll be the image in RP2. And we can take the intersection of that line with z equals 1 to see the image in this affine plane.

Now, notice that PGL3R, this matrix group, is eight dimensional. So we built PGL3R by equating all scalar multiples of a given matrix in GL3R. So GL3R was nine dimensional, but we lost a dimension when we went to PGL3R. And that's why I'm putting a 1 here, just to signify that that entry is not free. That matrix entry is not free.

But the rest are all free entries. They can range over all of the real numbers. So there's really one, two, three, four, five, six, seven, eight, eight dimensions of freedom in PGL3R. And let's try and understand these eight dimensions as eight one-dimensional families of transformations on the plane z equals 1, which we can actually visualize. So this gives us a way of wrapping our heads around this strange matrix group and its action.

So another thing to just notice here, so PGL3R may map certain elements, certain points, in z equals 1 to points that are not in z equals 1. Many points in z equals 1 will be mapped out of z equals 1 by an element of PGL3R. But there's a six dimensional subgroup, which I'll call AF for affine, which fixes the plane z equals 1. It just has zeros for these two entries. So here's an element of z equals 1.

When we multiply it by this matrix, whose entries here and here are 0, we'll get something in z equals 1. Because when you perform matrix multiplication, what do we get? We get a11x plus a12y plus b1. For the y-coordinate, we get a21x plus a22y plus b2. And finally, for the z-coordinate, we get 0 times x plus 0 times y plus 1 times 1, which is just 1. So the z-coordinate remains 1.

So matrices of this form, matrices in this subgroup AF, will fix the plane z equals 1. On the other hand, these two remaining dimensions of transformations, where these two coordinates are non-zero, those are still important. And they'll still give very, very interesting elements of PGL3R with very interesting actions. Except the only difference is that these may take certain points from this plane off of this plane. But that's fine.

They'll still take lines to lines in RP2 and be perfectly well-defined maps of RP2. And we can still visualize them via the affine chart, which takes those lines back onto this plane. It's just that it doesn't literally fix the plane. So the action of PGL3R on this affine plane is what I want to look at now. And what follows is just one way of breaking down the eight-dimensional group of transformations into one-dimensional subspaces.

So in general, I'm going to look at the image of xy1. This is the x-axis. This is the y-axis within the plane z equals 1. So over here, I have the x-axis, y-axis, and z-axis in R3. Here's my plane z equals 1.

I'm going to zoom in on z equals 1 now, look at a top-down view of it, and just focus on this x-axis and this y-axis within z equals 1. And look at eight different one-dimensional subspaces within PGL3R and how their actions look on this plane z equals 1. So we'll start with some transformation subgroups of PGL3R that you may have encountered, like the isometries and the affinities. And finally, we'll get to some new ones that we only really get to when we start looking at projective geometry and projection. So we'll start with a three-dimensional subgroup known as the isometries.

And within that, we'll start with the orientation-preserving isometries. So our first family I want to look at is the family of translation along the x-axis by a distance of b1. b1 is just a real number. So it's given by this matrix. The only non-zero, the only non-trivial entry here is this b1.

So this is a one-dimensional family. b1 can be any real number. And it corresponds to translation along the x-axis, like this. On the other hand, if we let b2 be non-zero, but make this entry 0, then we get translation along the y-axis by a distance of b2. And combining these, we get a general translation by a vector b1, b2.

And we can let those be any two real numbers and get translation by any vector we want. So a third family now that I want to bring up is also an orientation-preserving isometry. And this is counterclockwise rotation about the origin by an angle of theta. And it's given by this matrix. So here, I've put some non-trivial stuff into all four of these entries.

But really, there's just one parameter, this parameter theta. So this is also a one-dimensional family. It's just that I'm using all these entries to describe it. And here, you're seeing rotation by this angle theta. But you could use any other theta.

And altogether, that makes up the orientation-preserving isometries. That's a three-dimensional family. But now, let's expand to the full family of isometries, which include reflections. Those do not preserve orientation. So in particular, we have reflection about the y-axis, which is given by this matrix.

And notice this is just a single element. There's no free variables here. So there's no contribution to dimension. So the group of all isometries is also a three-dimensional space. And this is just flipping about reflection about the y-axis, as you can see.

So let's expand now beyond isometries and look at affinities, f. And these include certain transformations which distort the plane. So the fourth family I want to mention is stretching along the x-axis by a factor of lambda, also known as dilation. So here, it's given by this matrix, which has a non-trivial free entry in this upper left entry here. Lambda can be any real number, any non-zero real number, rather.

And it stretches along the x-axis by a factor of lambda. This is not an isometry. It doesn't preserve length or distance. But it does preserve other things.

So this is my fourth family. Very similarly, we can stretch along the y-axis by letting this entry of the matrix be non-trivial, non-zero, and keeping everything else 0 or 1. So this is, again, a one-dimensional family. And by using different values of mu, I can get different scaling factors and stretch along the y-axis by a factor of mu. So there's a sixth family. There's one more family of transformations in this affine group here.

And that's the shearing transformation, shearing along the x-axis by an angle of phi. So this involves putting a parameter in this entry of the matrix, tan of phi, where phi is the angle I'm shearing at. So this is a little bit like stretching. But rather than taking the square to a rectangle, I'm taking it to a parallelogram. I'm stretching this point, this upper right point, to any point I want along this in space.

So it's one parameter. I can choose how far I want to stretch it. So that is the shearing transformation along the x-axis. And together, all of these things I've described so far make up the family of the subgroup aff, the family of affinities. And up till now, we've been literally preserving that plane z equals 1.

We've been fixing that plane z equals 1. We're not actually taking any vectors in that plane off of that plane. But now we get to the really interesting things that are fully new. You wouldn't see these in your first course in linear algebra, probably, where we start allowing non-zero entries in these two coordinates here. Up till now, everything has had these as 0.

So let's move on to family 7. And this is in what we'll call the group of projectivities. And family 7 involves perspective distortion with respect to the x-axis. So it'll look like this. Perspective distortion with respect to the x-axis. Basically, you can imagine pushing the x-axis out away from us.

And it maps the line at infinity to the line x equals 1 over c1, where c1 is the entry in this lower left corner of the matrix. So you can very concretely realize this perspective distortion with a very simple matrix. And you can see that, yes, this is indeed the line at infinity. This point here is the point at infinity for all of these guys. Similarly, our eighth and final family involves a non-zero, non-trivial entry here, which can range freely over the real numbers, c2.

And it represents perspective distortion with respect to the y-axis, pushing the y-axis away from us, which looks like this, and maps the line at infinity to the line y equals 1 over c2. And you can see that the point at infinity of these guys is this point here. We can also put those together, these two perspective distortions, and get a general perspective distortion, sending the line at infinity to the line c1x plus c2y equals 1. We're using both c1 and c2 here, we can put entries in both of those, and combine them to get a more interesting distortion.

And we get a line at infinity here. So that gives us the eight families of transformations that make up our eight-dimensional space PGL3R. But this is just one way of imagining the eight families. For example, I actually arbitrarily made some arbitrary choices in how I presented it. I chose to showcase the shear transformation in the x direction. I could have also included the shear transformation in the y direction, and then excluded the rotation transformation.

That would have given one family that's explicitly linked to each of these entries, but it would have excluded rotation, which is such a natural transformation that it felt wrong to exclude it.

So we have a... really this is... there's no natural or correct way to break this up into eight families. It's up to you how you want to do it. So this is just a presentation that I personally liked.

So... But it hopefully gives you some intuition about how this group PGL3R acts on RP2, using this affine chart, c equals 1, to visualize. So going back to the synthetic analytic equivalence theorem, it states that projectivities from this plane z equals 1 to itself are in one-to-one correspondence with the elements of the matrix group PGL3R.

And the diff... So really everything we've been studying up till now can be framed in terms of PGL3R and RP2. But this new framework with PGL3R allows us to concretely apply any of the transformations involved in perspective shifts. So it's really crucial in, say, computer graphics or image correction. If you have a photograph of a document and you want to correct it, that involves a perspective transformation, and that can be captured by a matrix.

So it's very crucial in lots of practical applications. The proof of this theorem, this equivalence theorem, is a bit beyond the scope of this course, but only slightly. So if you're interested in checking it out, you can check out the book Foundations of Projective Geometry by Robin Hartshorn, and that gives one method of proving it. So in summary, when we shift perspective, there's eight dimensions worth of change that are happening. And there's really two approaches to understanding these changes. There's the synthetic approach involving a series of projectivities or projections from a plane to a plane to a plane back to itself, and that can map any four points to any other four points.

There's eight degrees of freedom, which we see in the four vertices of an initial square and a target square. Rather, an initial square can map to any target quadrilateral, and the four vertices of that quadrilateral, each of which have two dimensions, gives us eight degrees of freedom. On the other hand, the analytic approach uses a matrix in PGL3R, which acts on RP2, and that allows us to do numerical calculations. And again, we see eight degrees of freedom, but this time we see it in terms of the eight free matrix entries.