PATRICK WINSTON: Many of you, maybe most of you, will never have to work another search problem by hand in your natural life.

Others will want to take another run at it on the final.

I've been much criticized for this way of doing grading in the class.

But the way I look at it is that the relationship between students and instructors ought to be less adversarial than it used to be in the good old days when I was a student.

I especially remember an examination we took-- all of us took-- on Rayleigh scattering.

That was 803, I think.

And back in the old days, we had to take four physics courses, not just two.

And we had to take four math courses, not just two.

Back in the days when-- well, I was going to say men were men.

But most of us were men in those days.

I think we only had 20 women in our class.

Anyway, we had a quiz on Rayleigh scattering.

And I thought, well, this is pretty hard.

And then I got my quiz back.

26.

I thought, well, I've been found out.

I'm going to flunk out.

My father will make me go to law school.

I'll never attract anyone to marry.

Horrible things will happen.

Then the instructor announced the class average was 18.

I was two standard deviations above that.

They gave us the same exam two weeks later, and accounts vary.

Some people say that the class average went down.

Anyway, today we're going to study some stuff.

We're going to study some stuff that will make it possible for you to understand how you can do that computation in just a couple of seconds, even with the delays introduced by the redrawing.

Now this particular program-- I'm not real sure and I don't have a proof, but I think it will take more than the lifetime of the universe to find a legitimate coloring of the continental United States.

But by the end of the next class, you'll see how to do that lickety-split in just couple of seconds, even with the re-display delays.

Now we could, of course, do this in two ways.

One way is I could start off by saying, let C be a constraint satisfaction problem and just give you the result.

And anybody can do that.

That's great.

And you needed to learn some stuff that way.

But there are some things that you needed to learn, I think, a different way.

And that different way involves my telling you the story of how it all came to be.

This is a new field, pretty much, and therefore I know most of the people in it.

And I know all of the people who did the work that I'm going to tell you about today.

I'm going to tell you a little bit about the struggle for coming up with the ideas that led to one of the most powerful methods you'll learn about in the subject.

It all has to do, originally, with an attempt to make a computer capable of seeing.

And everybody said to themselves, well, let's start with seeing simple things, like children's blocks.

And so Adolfo Guzman, a graduate student of Marvin Minsky's, was charged with the summer project, which led to his Ph.D. thesis, of writing a program that could look at a line drawing and determine how many objects are in the line drawing.

So for example, there behind my outline is a line drawing.

And you believe instantly that there are two objects there, even though in some deep sense, it's ambiguous.

There are all sorts of ways that you could, through trickery, arrange something with even seven objects that look that way.

But most people would say there are two objects.

So Guzman set about the problem of figuring out how to do that.

And then his work was followed by Dave Huffman.

And his work was followed up by Dave Waltz.

And his work was followed up by Jane [? Froyter ?], who's not listed there quite yet.

And I want to tell you a little bit of story how that all happened.

So Guzman took a lot of pictures.

He went to Boston Baby, the precursor to Toys R Us, purchased a large block set on a government contract, and went about the business of taking a lot of pictures of them so he'd get a feel for the domain.

And eventually he decided that he could build a program that could determine that there are two objects here by using the lines as a quanta of evidence about which faces belong together in an object.

So he said, after studying these for a long time, that these drawings tend to have a lot of arrow-type junctions and a lot of fork-type junctions.

And when you see an arrow, it's a little bit of evidence that the objects on either side of the shaft are the same-- the face on either side of the shaft belong to the same object.

And over here, the fork-type junction suggests that pairwise, three pairs of faces seem to belong to the same object.

So with that idea, he could come back to a drawing like this and start decorating it with these so-called links, these quanta of evidence that faces belong to the same object.

And if I've done it right, you get something that looks like so.

It's a little hard to see what's going on on the drawing itself.

So let me number these.

Now I have an easier-to-deal-with picture.

There are two links between 1 and 2 and 1 and 3.

One link between 2 and 3.

One between 2 and 4.

Two here.

Two here.

And one each from all of these.

So Guzman produces this evidence for how the problem ought to be solved and then he begins to think about various ways of using the evidence.

So he could, and did, decide that one link is enough to presume that the faces belong to the same object.

And you can see that that's a little bit too liberal.

Because that says that the whole thing is just one object.

So Guzman rejected the one-link theory and went to the two-link theory.

And the two-link theory says, oh, well, let's see.

If we insist that two links should be there before we are willing to decide that it's the same object, then these three faces are pulled together into one object.

These three are pulled together into one object.

But alas, 7 doesn't share two links with anything, so it's left dangling out there.

So plainly, the two-link theory is too conservative.

So that as you would soon discover in any [INAUDIBLE] project, would lead you to a third theory, which is two lengths repeated.

So now that we've formed these super regions, we can come back and say we'll merge super regions if they're connected together by two or more links.

So nothing new happens up here.

But this super region is joined at 7 by two or more links.

So it pulls everything together like so.

So that worked fine.

Well, it didn't work fine.

There were lots of examples of situations where it didn't work-- in situations that were considered nonsensical by us humans because it seemed silly, the kind of conclusions that it reached.

So when Guzman presented this work in his Ph.D. thesis defense, it's said-- and who knows if it's true-- but it's said that Marvin Minsky thought it was great work and we should make him a professor.

It happens that Dave Huffman was also on the committee and said that it was ad hoc and the thesis should be rejected.

So we had two polar opposites of impressions there.

Dave Huffman, by the way-- you've heard that name before, I imagine.

It's the guy who invented Huffman coating as a term paper in a information theory course taught by Bob Fano.

He got an A, they say.

So Huffman didn't like it.

He thought it was a little bit too ad hoc.

It was too heuristic.

It didn't take advantage of anything we might know about physics.

And so people began to say, well, why does it work?

And when does it not work?

And that's just about the best question you can answer in a situation like this.

And by playing with some more of Guzman's pictures and reflecting on them, it turned out that it worked because the world is full of three-face junctions.

Or let me say three-face vertexes because they're out there in the world.

We'll reserve the word junction for something else.

And three-face vertexes generally project into either an arrow or a fork.

So Guzman's whole program worked on the weak backward conclusion that if you see one of those, it probably came from one of these.

So this is in the drawing.

That's in the world.

So by a process that's neither deduction or induction, but rather abduction, you see one of those guys.

And you say, well, they often come from three-face vertexes in the world, so if you see one, it probably came from a three-face vertex in the world.

That's abduction.

So once Huffman saw all that, being a mathematician, he began to think about how one might develop a different and better theory.

But we have to recognize that all this work started off with the efforts of Guzman, who was an experimentalist.

And Huffman was a mathematician.

So naturally, they approached the problem differently.

So Huffman says, I'm not going to concern myself too much with the actual problem that Guzman was trying to solve.

Rather, I'm going to work in a very simple world, which I can deal with mathematically.

So he decided that he was going to work in a world which had several characteristics.

Characteristic number one was that the world would be presented in general position.

That is to say, no screw cases.

So if you see a cube, it's going to look like this.

And it's not going to look like this.

So that's out.

And this is in.

And the idea here is that that's not general position, because if you perturb your point of view a little bit, you'll change those T junctions into forks and arrows.

So we know it had to deal with those kinds of weird kinds of cases.

So that's presumption number one.

Presumption number two is that we're going to be dealing with a world that's trihedral.

That means all vertexes out there are going to be formed from three planes.

So you're going to have the intersection of three planes.

So how many kinds of junctions can you see if the world is composed that way is the question that Huffman sent himself upon.

The next assumption was that there are going to be four kinds of lines.

See, Guzman only recognized two kinds of lines-- lines that have a length and lines that don't have a link.

And they don't have very much to do with the physical world.

So Huffman says, I want to get the constraint out of the physical world.

So I'm going to deal with four kinds of lines-- concave, convex, and boundary.

So each of those kinds of lines is going to have its own notation.

We'll call the convex lines plus, the concave lines minus, and a boundary lines are going to get an arrow on them, depending on which side of the object-- which side you would see the object if you walk along the direction of this little arrow.

STUDENT: Question.

Which is concave and which is convex?

You said something and you wrote the opposite.

PATRICK WINSTON: Yeah, sorry.

Thank you.

So concave, convex, and boundary.

Thank you.

So this down here, that's a concave line.

And that would get a minus label.

Oh, I don't know.

These lines you're seeing here, if there were stuff behind them instead of me, then all those would be convex lines.

Many times you see a boundary line.

For example, now you don't see both sides of that line down here at the bottom.

So that's a boundary line.

And the arrow would point in that direction, so as to keep the stuff of the object on the right as you walk along a kind of mathematical convention.

So three kinds of lines, four kinds of labels.

And there's some things left out.

That's because Huffman wanted to keep his problem simple, something he could manage by hand.

What's left out?

Cracks are left out.

Shadows are left out.

There's a presumption that there's nothing of interest here with respect to that.

So let's have a little vocabulary before I go on.

And I'll try to stick to it.

But there's the stuff out there, and that consists of vertexes and edges.

And over here, on the blackboard, we'll have junctions and lines.

And I'll try to stick to that vocabulary.

All right?

Yes, Christopher?

STUDENT: Didn't you say there were four types of lines?

PATRICK WINSTON: Yeah, there are four kinds of-- the question is, didn't I say there were four kinds of lines?

There are three kinds of lines, but since we can label a boundary line in either direction, we have four labels.

OK?

It depends on which side of the object the stuff is on.

And that will be clear, I think, as soon as I do an example.

So one, two, three assumptions.

A little bit of vocabulary.

So we have the possibility of making a complete catalog.

This is so simple.

We have the possibility of making a complete catalog of all the ways that lines can come together to form a junction with respect to these four labels.

Now at first you might say, oh my god, that will take a couple of years.

But maybe it won't take a couple of years.

And in the end, to perhaps your surprise, you discover that there are only 18 ways to arrange the labels around a junction in this world.

Everything else is excluded.

If you have something that's not in the set we're going to produce now, it can't be built with trihedral vertexes.

So I've listed up there six L's, five forks, four T's, and three arrows.

Let's see if I can figure out why there are those 18 and nothing else.

Well, if we have three vertexes coming together, that means there are eight octants, right?

And the stuff of the object may fill 1, 2, 3, 4, 5, 6, 7 or all eight octants.

Now of course, if you fill all eight octants, there's no junction.

So we don't consider that case.

If we don't fill any the octants, there's no junction.

There's no vertex.

So we don't consider that case.

But if just one of the eight octants is filled with stuff, then we can look at it from any of the seven remaining octants.

So right now, you're looking at it from one of the seven remaining octants.

And if I'm not mistaken, you're going to see a fork-style junction there, right?

And you're going to see a fork-style junction in which all of the edges are convex.

An unfortunate selection of names, because the linguists tell me that we index all of our words by the first forename.

And those have the same first forename, so they're easily confused.

So here's a fork-style junction.

And we know that one way that's legitimate for the lines to come in is with four pluses.

Now that's not the only way you can see that.

And here's another way.

There's an L-style junction, right?

And both of those are boundary lines.

And we want to draw the boundary line indicators on there so that if we walk along the arrows, the stuff of the object would be on the right.

So I suppose, to make it easier to me do my drawing, I should look at it this way.

And then I would say, well, that has to be a legitimate way of labeling a junction.

Are there any more?

Well, there's seven altogether, but many of them are the same.

There is one more that's a little different, though.

I can hold this guy up like so.

And if I'm holding at the right angle for you, you see an arrow-style junction, right?

Two boundaries, and the barb is convex.

So in this particular case, I've got that.

I've got that.

And I've got a plus there.

And that's there happen to be for the one octant filled with stuff case.

I happen to be able to reverse this, though.

And here's the seven octants filled case.

So that tells us that it's possible to have a vertex out there in a space that when reduced to a junction on the board deserves three minus labels, because all of these that you're seeing now are concave.

So another fork-style junction looks like so.

And since there's only one octant to look at from, that completes our analysis of the seven octants filled case.

Now we have a couple of other possibilities here.

We might have five octants filled with stuff.

So that means there are three octants that we can look from.

So let's suppose that these little chalk pieces are little people.

And they're looking at this junction down there, where this white object is fused with the table.

I'm fusing it with the table because I want to consider it to be one object.

We can view it from the three objects indicated by those three little chalk pieces.

And ask ourselves, well, in the event that we look at it from those three places, what do we see?

And if we look at it from the perspective of the piece of purple chalk-- I'll have to walk around and be sure-- looks like an arrow junction with two concaves and a convex.

Did I get that right?

So I'm looking at it from this perspective.

It's an arrow.

This is convex and these two are concave because I fused the paper box with the table.

Looking at it from the perspective of this blue guy-- let me rotate it so you can have a better understanding of the blue guy-- it looks like a concave line and a boundary.

So it's an L. And this one is a boundary.

And that's concave.

And by a kind of symmetry, we're also going to get that one from the other side, the third of the three octants.

Well, we're off and running.

But we still have an awful lot to go.

And we could manage to deal with it by thinking about this object once again.

But instead of this situation out here, to turn it around and look at this vertex.

Think about the junctions that it can produce.

I think I'll do that for you.

Because you really have to play with this and move it around a little bit to see how things are going to look.

So let me think about how that's going to work out.

I know that one of the possibilities is going to look like so.

I might as well not hide that from you.

It's going to be what happens when there's a little man looking up at the junction.

And this one's going to be minus.

And now we've got two more that are just like that.

Look like so.

And you say, oops.

You say, aren't those just a rotational variance of each other?

And the answer is sure.

I write them all down, because if you get a fork-style junction in space, there are three different ways it could be labeled.

Depending on which of the lines you put the minus label on.

So that takes care of that.

And then there's one more of these fork-style junctions.

And that's plus, plus, minus that derives from this case.

And there appear to be three more of these L-style junctions.

And they look like, let's see, plus, then plus.

I'm having to think this through as I go.

And then-- and that's it.

Well, what about the T's?

Well, in this world, the only way you can get a T is if some object is obscuring another object.

And if an object is obscuring a line, it can be any line at all.

So that's why the four remaining ones are easy beyond description to label.

And of course, the cross pieces of the T are all boundary lines, with the obscuring object on the right.

Now let's see.

We've taken care now of the one, three, five, and seven octants filled cases.

What about two, four, and six?

Well, it turns out you can have vertexes that are made that way too.

But they will have more than three faces.

They'll have six faces.

They'll be like what happens when an object comes together at a point, like so.

Like that.

You can play with it a little bit and see that if you have two, four, or six objects filled with stuff, there are more than three faces.

So we're going to ignore those.

So our constraint is going to be a little more severe than would be suggested by the terminology Huffman uses.

They're going to be trihedral all right, but they're also going to be three faces.

So we went to a lot of work there.

But what have we discovered?

We've discovered that in this world, this is a complete, 100% percent, nothing excluded, nothing else can be there, catalog of all possible ways the junctions can have line labels arranged around them.

There's nothing else in this world.

So that's a very powerful constraint.

So now let's see what we can do with it.

This example is usually more fun when the Red Sox are doing better, but they're not.

Yet we'll use it any way.

We're going to start with an object that looks like home plate.

And I'll ask you the question.

Can you build one of those?

I don't know.

Let's give it a shot.

We're going to assume that this object is hanging, floating in space.

So therefore, all of these lines around the boundary are boundary lines, like so.

Now that gets us off to a good-- it's just hanging in space, Christopher, all right?

You look confused.

It's just hanging in space so that all the lines around the edges are places where you see only one side has stuff on it.

So that enables us to just quickly run around the periphery and put arrow labels on all those outside lines.

Now we have a lot of arrow-style junctions on the boundary.

That's commonly the case.

So we can run over to our catalog of all possible labels, and we see that if we have an arrow with boundaries on its barbs, there's only one of those.

So I know instantly that there must be a plus on the shaft.

So we can come back here and take all these arrows here.

And label them with plus lines on their shafts.

Now a line can't change its nature along its length.

So if it's a plus line on one end, it's going to be a plus line on the other end.

All right?

So what else can we do?

Here deep inside is a fork-style junction.

It's got convex markers on both of those two lines.

So we go over to our catalog and say, what can we say about it, given that there are pluses on two of its lines?

Whoop, that means that the third one has to be a plus as well.

And now we're done.

We've labeled everything.

Except-- look at this.

What about that guy?

There's an L-style junction with pluses on both of its two lines.

Is there one of those in my catalog?

No.

Therefore, I haven't passed a necessary condition for constructability in the world that I've made.

You can't make it.

You can't construct it.

So Huffman's ideas give us a way of testing something to see if it's not possible for it to be in this world.

If it passes the test, does that mean it's possible?

No.

It's a necessary but not sufficient condition.

On this one-- blue-collar occupation-- on this one, maybe it will help if we put in another line.

For example, we could put a line like so.

You feel better about it now?

I don't know.

Let's see.

This has to be a plus for the same argument we used on several other arrows before.

That gives us an arrow-style junction here with a plus on everything.

Is there such a junction label?

No.

We lose.

It doesn't help.

You think you can make it, but you can't.

STUDENT: You could actually construct it as a 3-D object, though.

PATRICK WINSTON: He thinks you can construct it as a 3-D object.

Let me show you the next example, Christopher.

Consider this example.

Can you make that?

Your intuition is yes.

So let's label it.

Oh, I've already lost.

We just boost that up a little bit to make the situation more clear.

So already, I've got myself in a situation where I can't label that.

But you feel like you can make it.

So what's wrong?

What's wrong is-- what, Elliott?

STUDENT: You have an obscured-- or, we're presuming that we have an obscured [INAUDIBLE] alley from the upper-left corner to the [INAUDIBLE].

PATRICK WINSTON: Putting a little interpretation on what Elliott has said.

If you look at this situation back here, you get a four-faced junction there.

So you can make it.

But not with three faces.

So some of these look like you could make it.

But they can't be labeled because you need more than four faces at a junction.

And we can carry that idea back here.

You can make this OK.

But this junction, you've got two in the back and two here.

So it has four faces.

Same idea.

So that's Guzman's contribution.

That's Huffman's contribution.

Huffman was a mathematician.

But we wanted to build robots back in those days.

And neither one of these guys solved the problem of dealing with natural objects with shadows, with cracks, with more than trihedral vertexes in space, and what to do about that?

Well, that was a problem that another graduate student, David Waltz, set about to solve.

So Waltz decided that he would not be content unless he added cracks, shadows, non-trihedral vertexes.

Uh-- yeah, non-trihedral vertexes.

And light.

These considerations led Waltz to go from four labels to 50-plus.

Because he had to pack a lot of information into each of the labels he put on a line.

What kind of light was on the right?

What kind of light was on the left?

Maybe it's a crack.

Maybe it's a crack that-- all sorts of considerations.

Here we had 18 ways that lines can come together around junctions.

That went to thousands of junctions in Waltz's world.

So here's Guzman.

He writes a program that sort of works.

Down below, we have Huffman.

Huffman, who has a theory but solves the wrong problem.

So here comes Waltz, and he's trying to solve the right problem with a satisfying theory that has a generalizable principle.

So when we get all through this, we'll talk about criteria for success.

And we'll conclude that to have a really successful thing, you need a problem, to start with.

You need a method that works.

And you have to show that it works because of some principal.

So Guzman had the problem and something that worked.

Huffman had a method which worked on the wrong problem.

And it's left to Waltz to bring it all together.

So Waltz does all this work.

And now he has, instead at 18 labels, he has thousands.

Instead of four-line labels, he has more than 50.

So naturally, it becomes difficult to work these by hand.

We were able to work those Huffman examples by hand.

We started with labeling everything on the boundary and worked our way in.

There's no particular method there.

It was just easy to work out the puzzle.

But poor Waltz, he didn't have that luxury.

So he might have, in a typical scene, he might have tens or even hundreds of junctions to label and no easy way of dealing with it.

So naturally writes himself a depth-first search program.

So here is vertex, or rather junction number 1.

There are many choices for which label can be used on it.

And for each of those choices, whatever he's decided junction number 2 is has its own suite of possibilities.

And so it becomes a simple depth-first search problem, right?

So in actuality, as soon as Waltz-- he was my office mate at the time.

I can tell you this for a fact.

As soon as he wrote this program, he kept looking over at the computer-- they were big in those days.

They all had lights.

So you looked over at the computer to see if the lights were still blinking.

Because he'd start this depth-first search program up and nothing would happen.

He thought the computer had crashed.

Nothing happened.

Why did nothing happen?

Because the search base is exponential and much too big for an ordinary computer, or maybe even an ordinary universe.

So Waltz has to do something else.

He has to come up with a new method for using all these labels that he's-- after about a year and a half's worth of hard work, with lots of paper on his desk in little blocks.

After year and a half of hard work getting all these junction labels figured out, he then has to come up with a method for figuring out how to use them.

And so we don't know whether to think his biggest contribution was that label set or his method.

And probably both deserve about equal billing.

Oh, I don't know how to explain what Waltz did.

Well, one way is to do an example.

And I think I will hazard an example.

Let's see.

Let me find some space.

I think I'm reduced to going over here.

But that will be convenient, since the line labels are here.

Here's my example.

And you say, how can I give you just part of a picture?

Well, you can assume I'm looking at this through a window.

So the edge of the window form boundary lines, and they exert no constraint whatsoever on what's behind them.

So this is a legitimate drawing to have to think about.

By the way, is this ambiguous?

Or do you get a firm sense that there's a unique interpretation of all those lines?

I think there's a unique interpretation of all those lines.

What I'm going to do is I'm actually going to solve this problem using Huffman's labels but Waltz's method.

Because I can't simulate on the blackboard something with 50 line types and thousands of line junctions.

So I'm going to use Huffman's set to demonstrate Waltz's algorithm.

So Waltz's algorithm involves starting out by plopping on some junction all of the possible labels that the answer has to be drawn from.

So let me number these in the order that we're going to visit them.

Like so.

And so far, I've just put down the three fork options that are resident on that first junction.

And I have to take note of exactly what they do with the lines that come out of the junction.

So let's see.

I'll just copy them down.

One possibility is this one.

Another possibility is this one.

And another possibility is plus, plus, minus.

Oops, I've got plus, plus.

No, that's right.

So that right so far?

All I've done is copy the junction labelings from my library.

And at this point, Waltz's algorithm says there's nothing else to do but go on to junction number two.

And unfortunately, sadly, there are lots of labelings that have to be considered on junction number 2.

Six of them.

1, 2, 3, 4, 5, 6.

So one of them looks like that.

Another one looks like that.

One of them is plus here, arrow in.

Another one is plus here, arrow out.

Another one is minus here, arrow down.

And minus here, arrow up.

I think I've copied those all right.

But now, having copied those down, Waltz's algorithm looks around at the neighboring junctions and says, are any of the things that I just placed on junction two disallowed by what I've already placed on a neighboring junction?

So it looks over here in step number two.

And it sees that these three arrows require the line that joins junctions 1 and 2 to be either minus or plus.

So of the six possibilities, I can only keep the ones that are likewise content to put a plus on that line that joins the two.

So that means that the influence flowing from junction 1 eliminates that one, eliminates that one, eliminates that one, and eliminates that one.

So half of them are gone.

All the ones that try to put a boundary line on that line between 1 and 2 are disallowed.

All right?

Now likewise, we could say, well, of the remaining ones, do they restrict what I can do at junction 1?

So let's see.

Here's a minus.

And here's a plus.

So all these possibilities over here are still alive.

So now, continuing on, we have to see what we can do at junction 3.

These are arrow labels again.

So we have to copy exactly the same labels set as we had before.

And now we look down at junction 2 and say, well, what does that tell me about the three that I've just placed?

If we look up from junction 2 to see what kind of constraints it puts on here, we have this one alive and this one alive.

I guess we've eliminated four of the six.

So we have these two alive.

And they both but boundary lines-- I think I must have had this boundary line wrong, right?

No, that's right.

Oh yes, I see.

Plus.

This one goes-- hang on a second.

You let me do something wrong.

So plus is out.

And that must be one that goes-- this minus goes up.

Oh, yes.

I'm too hasty and uncertain about what I was doing.

So let's see.

This guy has a boundary going down.

And this guy has a boundary going up.

All of the others have been eliminated.

So that means that something that tries to put a concave line there is gone.

And something that tries to put a plus line there is gone.

So the influence flowing up in this direction in the third step eliminates that guy and eliminates that guy, leaving only this guy.

But now, the thing I was worried about is you have to also at this point go down to 2 and see if there's any further constraint on what can survive down there, based on what has happened over here at junction 3.

Now let's see.

This one goes up, which is compatible with a survivor.

But this one goes down, which is not compatible with a survivor.

So when I bring this down in step three, this guy is eliminated.

So now I'm down to just one interpretation for what can be going on at vertex number 2.

And one interpretation for vertex number 3.

Now let's see.

This can propagate.

So now that I've made a change on vertex number 2, I have to also see if that causes a change at vertex number 1.

So it's propagating through.

And now I can see that the only possibility here is a minus, the label that's coming down from our survivor.

So that eliminates these two.

Whew!

It's hard to do this by hand, but I've got three of the four things labeled.

And even with just three of the four labeled, I'm down to a single interpretation for all of the junctions and the lines between them.

So there's one left.

We have to deal with that fork vertex.

We better deal with it, because for all we know, this is not a legal drawing in this world.

We have five fork vertexes to place.

But you know what?

I don't have to draw much here, because I know this is forced to be a plus now.

And this is forced to be a plus now.

And there's only one fork vertex with any pluses on it at all.

So now I can come through and say, well, the only possible survivor is this one.

These are gone.

And now I have an interpretation for all of the junctions.

And I see that the winners are this one.

And this one.

And this one.

And this one.

So I've got a unique interpretation.

This line is convex.

This one is concave.

This is a boundary.

That's a boundary.

And this line is convex, and that's convex.

Now that's a lot of work.

So I better check and make sure I got it right.

You'd like to see this demonstrated to make sure I haven't made a mistake.

I'm sure of that.

Let's see.

That it?

So each of the places where a line is obscured has four possibilities, labeled E. The arrow junctions are labeled A.

The forks-- there are five of them-- at the fork junction 5.

So let's just step through here and see what happens.

Boom.

I've got it.

I did do it right.

So let's try some more.

What do you think will happen with this one?

Unique solution?

It stopped.

Bug in my program?

Unthinkable.

What's happened?

It is genuinely ambiguous.

It can be something hanging down from the ceiling.

Or it could be something that we can think of as a step going up from left to right.

Let's look at something more complicated.

You think it'll work?

Not enough constraint for us to figure that one out.

It's equally ambiguous, but a little bit larger example.

What about this one?

Yeah, but the stuff is creeping up from the lower left up to the upper right.

Yeah, bingo.

It worked.

It's unambiguous.

It's variation on the same theme we had before.

But let me, just for fun, take these two lines out.

What do you think will happen now?

Seems to be doing just fine until it hits the upper right-hand corner and discovers it can't label stuff.

So it propagates back down.

And what looked OK in the lower left is no good after all.

So these results are kind of consistent with what we humans do when we look at these kinds of things.

So it's very likely that we, in our heads, do have some constraint propagation apparatus that we use in vision.

But putting that aside, we can think about other kinds of intelligence different from human, that might use this kind of mechanism to solve problems that involve a lot of constraint in finding a solution.

So here, we saw the constraint propagation activity at work on line drawing analysis.

But next time, what we're going do is we're going to see at work in map coloring.

And who cares about map coloring?

People who do scheduling, because that turns out to be the same problem.