

MITOCW | Lec-01

PATRICK WINSTON: Welcome to 6034.

I don't know if I can deal with this microphone.

We'll see what happens.

It's going to be a good year.

We've got [INAUDIBLE] a bunch of interesting people.

It's always interesting to see what people named their children two decades ago.

And I find they were overwhelmed with Emilys.

And there are not too many Peters, Pauls, and Marys, but enough to call forth a suitable song at some point.

We have lots of Jesses of both genders.

We have a [INAUDIBLE] of both genders.

And we have a Duncan, where's Duncan?

There you are, Duncan.

You've changed your hairstyle.

I want to assure use that the Thane of Cawdor is not taking the course this semester.

What I'm going to do is tell you about artificial intelligence today, and what this subject is about.

There's been about a 10% percent turnover in the roster in the last 24 hours.

I expect another 10% turnover in the next 24 hours, too.

So I know many of you are sightseers, wanting to know if this is something you want to do.

So I'm going to tell you about what we're going to do this semester, and what you'll know when you get out of here.

I'm going to walk you through this outline.

I'm going to start by talking about what artificial intelligence is, and why we do it.

And then I'll give you a little bit of the history of artificial intelligence, and conclude with some of the covenants by which we run the course.

One of which is no laptops, please.

I'll explain why we have these covenants at the end.

So what is it?

Well, it must have something to do with thinking.

So let's start up here, a definition of artificial intelligence, by saying that it's about thinking, whatever that is.

My definition of artificial intelligence has to be rather broad.

So we're going to say it's not only about thinking.

It's also about perception, and it's about action.

And if this were a philosophy class, then I'd stop right there and just say, in this subject we're going to talk about problems involving thinking, perception, and action.

But this is not a philosophy class.

This a Course six class.

It's an engineering school class.

It's an MIT class.

So we need more than that.

And therefore we're going to talk about models that are targeted at thinking, perception, and action.

And this should not be strange to you, because model making is what MIT is about.

You run into someone at a bar, or relative asks you what you do at MIT, the right knee jerk reaction is to say, we learned how to build models.

That's what we do at MIT.

We build the models using differential equations.

We build models using probabilities.

We build models using physical and computational simulations.

Whatever we do, we build models.

Even in humanities class, MIT approach is to make models that we can use to explain the past, predict the future, understand the subject, and control the world.

That's what MIT is about.

And that's what this subject is about, too.

And now, our models are models of thinking.

So you might say, if I take this classic will I get smarter?

And the answer is yes.

You will get smarter.

Because you'll have better models of your own thinking, not just the subject matter of the subject, but better models of your own thinking.

So models targeted at thinking, perception, and action.

We know that's not quite enough, because in order to have a model, you have to have representation.

So let's say that artificial intelligence is about representations that support the making of models to facilitate an understanding of thinking, perception, and action.

Now you might say to me, well what's a representation?

And what good can it do?

So I'd like to take a brief moment to tell you about gyroscopes.

Many of you have friends in mechanical engineering.

One of the best ways embarrass them is to say here's a bicycle wheel.

And if I spin it, and blow hard on it right here, on the edge of the wheel, is going to turn over this way or this way?

I guarantee that what they will do is they'll put their hand in an arthritic posture called the right hand screw rule, aptly named because people who use it tend to get the right answer about 50% of the time.

But we're never going to make that mistake again.

Because we're electrical engineers, not mechanical engineers.

And we know about representation.

What we're going to do is we're going to think about it a little bit.

And we're going to use some duct tape to help us think about just one piece of the wheel.

So I want you to just think about that piece of the wheel as the wheel comes flying over the top, and I blow on it like that.

What's going to happen to that one piece?

It's going to go off that way, right?

And the next piece is going to go off that way too.

So when it comes over, it has to go that way.

Let me do some ground f here just to be sure.

It's very powerful feeling.

Try it.

We need a demonstration.

I don't anybody think that I'm cheating, here.

So let's just twist it one way or the other.

So that's powerful pull, isn't it.

Alex is now never going to get the gyroscope wrong, because he's got the right representation.

So much of what you're going to accumulate in this subject is a suite of representations that will help you to build programs that are intelligent.

But I want to give you a second example, one a little bit more computational.

But one of which was very familiar to you by the time you went to first grade, in most cases.

It's the problem of the farmer, the fox, the goose, and the grain.

There's a river, a leaky rowboat that can only carry the farmer, and one of his four possessions.

So what's the right representation for this problem?

It might be a picture of the farmer.

It might be a poem about the situation, perhaps a haiku.

We know that those are not the right representation.

Somehow, we get the sense that the right representation most involve something about the location of the participants in this scenario.

So we might draw a picture that looks like this.

There's the scenario, and here in glorious green, representing our algae infested rivers is the river.

And here's the farmer, the fox, the goose, and the grain.

An initial situation.

Now there are other situations like this one, for example.

We have the river, and the farmer, and the goose is on that side.

And the fox and the grain is on that side.

And we know that the farmer can execute a movement from one situation to another.

So now we're getting somewhere where with the problem.

This is at MIT approach to the farmer, fox, goose, and grain problem.

It might have stumped you when you were a little kid.

How many such situations are there?

What do you think, Tanya?

It looks to me like all four of individuals can be on one side or the other.

So for every position the farmer can be, each of the other things can be on either side of the river.

So it would be two to the fourth she says aggressively and without hesitation.

Yes, two to the fourth, 16 possibilities.

So we could actually draw out the entire graph.

It's small enough.

There's another position over here with the farmer, fox, goose, and grain.

And in fact that's the one we want.

And if we draw out the entire graph, it looks like this.

This is a graph of the situations and the allowed connections between them.

Why are there not 16?

Because the other-- how many have I got?

Four?

10?

The others are situations in which somebody gets eaten.

So we don't want to go to any of those places.

So having got the representation, something magical has happened.

We've got our constraints exposed.

And that's why we build representations.

That's why you algebra in high school, because algebraic notation exposes the constraints that make it possible to actually figure out how many customers you get for the number of advertisements you place in the newspaper.

So artificial intelligence is about constraints exposed by representations that support models targeted to thinking-- actually there's one more thing, too.

Not quite done.

Because after all, in the end, we have to build programs.

So it's about algorithms enabled by constraints exposed by representations that model targeted thinking, perception, and action.

So these algorithms, or we might call them just as well procedures, or we might call them just as well methods, whatever you like.

These are the stuff of what artificial intelligence is about-- methods, algorithms, representations.

I'd like to give you one more example.

It's something we call, in artificial intelligence, generated test.

And it's such a simple idea, you'll never hear it again in this subject.

But it's an idea you need to add to your repertoire of problem solving methods, techniques, procedures, and algorithms.

So here's how it works.

Maybe I can explain to best by starting off with an example.

Here's a tree leaf I picked off a tree on the way over to class.

I hope it's not the last of the species.

What is it, what kind of tree?

I don't know.

I never did learn my trees, or my colors, or my multiplication tables.

So I have to go back to this book, the Audubon Society Field Guide to North American Trees.

And how would I solve the problem?

It's pretty simple.

I just turn the pages one at a time, until I find something that looks like this leaf.

And then I discover it's a sycamore, or something.

MIT's full of them.

So when I do that, I do something very intuitive, very natural, something you do all the time.

But we're going to give it a name.

We're going to call it generate and test.

And generate and test method consists of generating some possible solutions, feeding them into a box that tests them, and then out the other side comes mostly failures.

But every once in a while we get something that succeeds and pleases us.

That's what I did with the leaf.

But now you have a name for it.

Once you have a name for something, you get power over it.

You can start to talk about it.

So I can say, if you're doing a generate and test approach to a problem, you better build a generator with certain properties that make generators good.

For example, they should not be redundant.

They shouldn't give you the same solution twice.

They should be informable.

They should be able to absorb information such as, this is a deciduous tree.

Don't bother looking at the conifers.

So once you have a name for something, you can start talking about.

And that vocabulary gives you power.

So we call this the Rumpelstiltskin Principle perhaps The first of our powerful ideas for the day.

This subject is full of powerful ideas.

There will be some in every class.

Rumpelstiltskin Principle says that once you can name something, you get power over it.

You know what that little thing is on the end of your shoelace?

It's interesting.

She's gesturing like mad.

That's something we'll talk about later, too-- motor stuff, and how it helps us think.

What is it?

No one knows?

It's an ag something, right?

It's an aglet, very good.

So once you have the name, you can start to talk about.

You can say the purpose of an aglet is pretty much like the whipping on the end of a rope.

It keeps the thing from unwinding.

Now you have a place to hang that knowledge.

So we're talking about this frequently from now into the rest of the semester, the power of being able to name things.

Symbolic labels give us power over concepts.

While we're here I should also say that this is a very simple idea, generate and test.

And you might be tempted to say to someone, we learned about generate and test today.

But it's a trivial idea.

The word trivial is a word I would like you to purge from your vocabulary, because it's a very dangerous label.

The reason it's dangerous is because there's a difference between trivial and simple.

What is it?

What's the difference between labeling something as trivial and calling it simple?

Yes?

Exactly so.

He says that simple can be powerful, and trivial makes it sound like it's not only simple, but of little worth.

So many MIT people miss opportunities, because they have a tendency to think that ideas aren't important unless they're complicated.

But the most simple ideas in artificial intelligence are often the most powerful.

We could teach an artificial intelligence course to you that would be so full of mathematics it would make a Course 18 professor gag.

But those ideas would be merely gratuitously complicated, and gratuitously mathematical, and gratuitously not simple.

Simple ideas are often the most powerful.

So where are we so far?

We talked about the definition.

We talked about an example of a method.

Showed you a representation, and perhaps also talked about the first idea, too.

You've got the representation right, you're often almost done.

Because with this representation, they can immediately see that there are just two solutions to this problem, something that wouldn't have occurred to us when we were little kids, and didn't think to draw the [? state ?] diagram.

There's still one more thing.

In the past, and in other places, artificial intelligence is often taught as purely about reasoning.

But we solve problems with our eyes, as well as our symbolic apparatus.

And you solved that problem with your eyes.

So I like to reinforce that by giving you a little puzzle.

Let's see, who's here?

I don't see [? Kambe, ?] but I'll bet he's from Africa.

Is anyone from Africa?

No one's from Africa?

No?

Well so much the better-- because they would know the answer to the puzzle.

Here's the puzzle.

How many countries in Africa does the Equator cross?

Would anybody be willing to stake their life on their answer?

Probably not.

Well, now let me repeat the question.

How many countries in Africa does the Equator cross?

Yeah, six.

What happened is a miracle.

The miracle is that I have communicated with you through language, and your language system commanded your visual system to execute a program that involves scanning across that line, counting as you go.

And then your vision system came back to your language system and said, six.

And that is a miracle.

And without understanding that miracle, we'll never have a full understanding of the nature of intelligence.

But that kind of problem solving is the kind of problem solving I wish we could teach you a lot about it.

But we can't teach you about stuff we don't understand.

We [INAUDIBLE] for that.

That's a little bit about the definition and some examples.

What's it for?

We can deal with that very quickly.

If we're engineers, it's for building smarter programs.

It's about building a tool kit of representations and methods that make it possible to build smarter programs.

And you will find, these days, that you can't build a big system without having embedded in it somewhere the ideas that we talk about in the subject.

If you're a scientist, there's a somewhat different motivation.

But it amounts to studying the same sorts of things.

If you're a scientist, you're interested in what it is that enables us to build a computational account of intelligence.

That's the part that I do.

But most this subject is going to be about the other part, the part that makes it possible for you to build smarter programs.

And some of it will be about what it is that makes us different from the chimpanzees with whom we share an enormous fraction of our DNA.

It used to be thought that we share 95% of our DNA with chimpanzees.

Then it went up to 98.

Thank God it stopped about there.

Then it actually went back a little bit.

I think we're back down to 94.

How about if we talk a little bit now about the history of AI, so we can see how we got to where we are today?

This will also be a history of AI that tells you a little bit about what you'll learn in this course.

It all started with Lady Lovelace, the world's first programmer, who wrote programs about 100 years before there were computers to run them.

But it's interesting that even in 1842, people were hassling her about whether computers could get really smart.

And she said, "The analytical engine has no pretensions to originate anything.

It can do whatever we know how to order it to perform." Screwball idea that persists to this day.

Nevertheless, that was the origin of it all.

That was the beginning of the discussions.

And then nothing much happened until about 1950, when Alan Turing wrote his famous paper, which introduced the Turing test.

Of course, Alan Turing had previously won the Second World War by breaking the German code, the Ultra Code, for which the British government rewarded him by driving him to suicide, because he happened to be homosexual.

But Turing wrote his paper in 1950, and that was the first milestone after Lady Lovelace's comment in 1842.

And then the modern era really began with a paper written by Marvin Minsky in 1960, titled "Steps Toward Artificial Intelligence." And it wasn't a long after that Jim [? Slagle, ?] a nearly blind graduate student, wrote a program that did symbolic integration.

Not adding up area under a curve, but doing symbolic integration just like you learn to do in high school when you're a freshman.

Now on Monday, we're going to talk about this program.

And you're going to understand exactly how it works.

And you can write one yourself.

And we're going to reach way back in time to look at that program because, in one day discussing it, talking about it, will be in itself a miniature artificial intelligence course.

Because it's so rich with important ideas.

So that's the dawn age, early dawn age.

This was the age of speculation, and this was the dawn age in here.

So in that early dawn age, the integration program took the world by storm.

Because not everybody knows how to do integration.

And someone, everyone, thought that if we can do integration today, the rest of intelligence will be figured out tomorrow.

Too bad for our side it didn't work out that way.

Here's another dawn age program, the Eliza [? thing ?].

But I imagine you'd prefer a demonstration to just reading it, right?

Do you prefer a demonstration?

Let's see if we can demonstrate it.

This is left over from a hamentashen debate of a couple of years ago.

How do you spell hamentashen, anybody know?

I sure hope that's right.

It doesn't matter.

Something interesting will come.

OK, your choice.

Teal?

Burton House?

Teal.

So that's dawn age AI.

And no one ever took that stuff seriously, except that it was a fun [INAUDIBLE] project level thing to work out some matching programs, and so on.

The integration program was serious.

This one wasn't.

This was serious, programs that do geometric analogy, problems of the kind you find on intelligence tests.

Do you have the answer to this?

A is to B as C is to what?

That would be 2, I guess.

What's the second best answer?

And the theories of the program that solve these problems are pretty much identical to what you just figured out.

In the first case you deleted the inside figure.

And the second case is, the reason you got four is because you deleted the outside part and grew the inside part.

There's another one.

I think this was the hardest one it got, or the easiest one it didn't get.

I've forgotten.

A is to B as C is to 3.

In the late dawn age, we began to turn our attention from purely symbolic reasoning to thinking a little bit about perceptual apparatus.

And programs were written that could figure out the nature of shapes and forms, such as that.

And it's interesting that those programs had the same kind of difficulty with this that you do.

Because now, having deleted all the edges, everything becomes ambiguous.

And it may be a series of platforms, or it may be a series of-- can you see the saw blade sticking up if you go through the reversal?

Programs were written that could learn from a small number of examples.

Many people think of computer learning as involving leading some neural net to submission with thousands of trials.

Programs were written in the early dawn age that learned that an arch is something that has to have the flat part on top, and the two sides can't touch, and the top may or may not be a wedge.

In the late dawn age, though, the most important thing, perhaps, was what you look at with me on Wednesday next.

It's a rule-based expert systems.

And a program was written at Stanford that did diagnosis of bacterial infections of the blood.

It turned out to do it better than most doctors, most general practitioners.

It was never used, curiously enough.

Because nobody cares what your problem actually is.

They just give you a broad spectrum antibiotic that'll kill everything.

But this late dawn age system, the so-called [INAUDIBLE] system, was the system that launched a thousand companies, because people started building expert systems built on that technology.

Here's that you don't know you used, or that was used on your behalf.

If you go through, for example, the Atlanta airport, your airplane is parked by a rule-based expert system that knows how to park aircraft effectively.

It saves Delta Airlines about to \$0.5 million a day of jet fuel by being all smarter about how to park them.

So that's an example of an expert system that does a little bit of good for a lot of people.

There's Deep Blue.

That takes us to the next stage beyond the age of expert systems, and the business age.

It takes us into this age here, which I call the bulldozer age, because this is the time when people began to see that we had at our disposal unlimited amounts of computing.

And frequently you can substitute computing for intelligence.

So no one would say that Deep Blue does anything like what a human chess master does.

But nevertheless, Deep Blue, by processing data like a bulldozer processes gravel, was able to beat the world champion.

So what's the right way?

That's the age we're in right now.

I will of course be inducing programs for those ages as we go through the subject.

There is a question of what age we're in right now.

And it's always dangerous to name an age when you're in it, I guess.

I like to call it the age of the right way.

And this is an age when we begin to realize that that definition up there is actually a little incomplete, because much of our intelligence has to do not with thinking, perception, and action acting separately, but with loops that tie all those together.

We had one example with Africa.

Here's another example drawn from a program that has been under development, and continues to be, in my laboratory.

We're going to ask the system to imagine something.

SYSTEM: OK.

I will imagine that a ball falls into a bowl.

OK.

I will imagine that a man runs into a woman.

PATRICK WINSTON: You see, it does the best that it can if it doesn't have a good memory of what these situations actually involve.

But having imagined the scene it can then-- SYSTEM: Yes.

I have learned from experience that contact between a man and a woman appeared because a man runs into a woman.

PATRICK WINSTON: Having imagined the scene, it can then read the answers using its visual apparatus on the scene that it imagined.

So just like what you did with Africa, only now it's working with its own visual memory, using visual programs.

SYSTEM: OK.

I will imagine that a man gives a ball to a man.

PATRICK WINSTON: I know this looks like slugs, but they're actually distinguished professors.

It always does the best it can.

SYSTEM: OK.

I will imagine that a man flies.

PATRICK WINSTON: It's the best that it can do.

So that concludes our discussion of the history.

And I've provided you with a little bit of a glimpse of what we're going to look at as the semester unfolds.

Yes, Chris?

CHRIS: Is it actually a demonstration of something?

Does it have a large database of videos?

PATRICK WINSTON: No, it has a small database videos.

CHRIS: But it's intelligently picking among them based on-- PATRICK WINSTON: Based on their content.

So if you say imagine that a student gave a ball to another student, it imagines that.

You say, now does the other student have the ball?

Does the other student take the ball?

It can answer those questions because it can review the same video and see the take as well as the give in the same video.

So now we have to think about why we ought to be optimistic about the future.

Because we've had a long history here, and we haven't solved the problem.

But one reason why we can feel optimistic about future is because all of our friends have been on the march.

And our friends include the cognitive psychologists, the [? developmental ?] psychologists, the linguists, sometimes the philosophers, and especially the paleoanthropologists.

Because it is becoming increasingly clear why we're actually different from the chimpanzees, and how we got to be that way.

The high school idea is that we evolved through slow, gradual, and continuous improvement.

But that doesn't seem to be the way it happened.

There are some characteristics of our species that are informative when it comes to guiding the activities of people like me.

And here's what the story seems to be from the fossil record.

First of all, we humans have been around for maybe 200,000 years in our present anatomical form.

If someone walked through the door right now from 200,000 years ago, I imagine they would be dirty, but other than that-- probably naked, too-- other than that, you wouldn't be able to tell the difference, especially at MIT.

And so the ensuing 150,000 years was a period in which we humans didn't actually amount to much.

But somehow, shortly before 50,000 years ago, some small group of us developed a capability that separated us from all other species.

It was an accident of evolution.

And these accidents may or may not happen, but it happened to produce us.

It's also the case that we probably necked down as a species to a few thousand, or maybe even a few hundred individuals, something which made these accidental changes, accidental evolutionary products, more capable of sticking.

This leads us to speculate on what it was that happened 50,000 years ago.

And paleoanthropologists, Noam Chomsky, a lot of people reached similar conclusions.

And that conclusion is-- I'll quote Chomsky.

He's the voice of authority.

"It seems that shortly before 50,000 years ago, some small group of us acquired the ability to take two concepts, and combine them to make a third concept, without disturbing the original two concepts, without limit." And from a perspective of an AI person like me, what Chomsky seems to be saying is, we learned how to begin to describe things, in a way that was intimately connected with language.

And that, in the end, is what separates us from the chimpanzees.

So you might say, well let's just study language.

No, you can't do that, because we think with our eyes.

So language does two things.

Number one, it enables us to make descriptions.

Descriptions enable us to tell stories.

And storytelling and story understanding is what all of education is about.

That's going up.

And going down enables us to marshal the resources of our perceptual systems, and even command our perceptual systems to imagine things we've never seen.

So here's an example.

Imagine running down the street with a full bucket of water.

What happens?

Your leg gets wet.

The water sloshes out.

You'll never find that fact anywhere on the web.

You've probably never been told that that's what happens when you run down the street with a full bucket of water.

But you easily imagine this scenario, and you know what's going to happen.

There was internal imagination simulation.

We're never going to understand human intelligence until we can understand that.

Here's another example.

Imagine running down the street with a full bucket of nickels?

What happens?

Nickels weigh a lot.

You're going to be bent over.

You're going to stagger.

But nobody ever told you that.

You won't find it anywhere on the web.

So language is at the center of things because it enables storytelling going up, and marshalling the resources of the perceptual apparatus, going down.

And that's where we're going to finish the subject the semester, by trying to understand more about that phenomenon.

So that concludes everything I wanted to say about the material and the subject.

Now I want to turn my attention a little bit to how we are going to operate the subject.

Because there are many characteristics of the subject that are confusing.

First of all, we have four kinds of activities in the course.

And each of these has a different purpose.

So I did the lectures.

And the lectures are supposed to be an hour about introducing the material and the big picture.

They're about powerful ideas.

They're about the experience side of the course.

Let me step aside and make a remark.

MIT is about two things.

It's about skill building, and it's about big ideas.

So you can build a skill at home, or at Dartmouth, or at Harvard, or Princeton, or all those kinds of places.

But the experience you can only get at MIT.

I know everybody there is to know in artificial intelligence.

I can tell you about how they think.

I can tell you about how I think.

And that's something you're not going to get any other place.

So that's my role, as I see it, in giving these lectures.

Recitations are for buttressing and expanding on the material, and providing a venue that's small enough for discussion.

Mega recitations are [? a usual ?] components of the course.

They're taught at the same hour on Fridays.

Mark Seifter, my graduate student, will be teaching those.

And those are wrapped around past quiz problems.

And Mark will show you how to work them.

It's very important component to the subject.

And finally the tutorials are about helping you with the homework.

So you might say to me, well, do I really need to go to class?

I like to say that the answer is, only if you like to pass the subject.

But you are MIT students.

And MIT people always like to look at the data.

So this is a scattergram we made after the subject was taught last fall, which shows the relationship between attendance at lectures and the grades awarded in the course.

And if you're not sure what that all means, here's the regression line.

So that information is a little suspect for two reasons, one of which is we asked people to self report on how many lectures they thought they attended.

And our mechanism for assigning these numerical grades is a little weird.

And there's a third thing, too, and that is, one must never confuse correlation with cause.

You can think of other explanations for why that trend line goes up, different from whether it has something to do with lectures producing good grades.

You might ask how I feel about the people up there on the other upper left hand corner.

There are one or two people who were near the top of the subject who didn't go to class at all.

And I have mixed feelings about that.

You're adults.

It's your call.

On the other hand, I wish that if that's what you do habitually in all the subjects you take at MIT, that you would resign and go somewhere else, and let somebody else take their slot.

Because you're not benefiting from the powerful ideas, and the other kinds of things that involve interaction with faculty.

So it can be done.

But I don't recommend it.

By the way, all of the four activities that we have here show similar regression lines.

But what about that five point scale?

Let me explain how that works to you.

We love to have people ask us what the class average is on a quiz.

Because that's when we get to use our blank stare.

Because we have no idea what the class average ever is on any quiz.

Here's what we do.

Like everybody else, we start off with a score from zero to 100.

But then we say to ourselves, what score would you get if you had a thorough understanding of the material?

And we say, well, for this particular exam, it's this number right here.

And what score would you get if you had a good understanding of the material?

That's that score.

And what happens if you're down here is that you're following off the edge of the range in which we think you need to do more work.

So what we do is, we say that if you're in this range here-- following MIT convention with GPAs and stuff, that gets you a five.

If you're in this range down here, there's a sharp drop off to four.

If you're in this range down here, there's a sharp fall off to three.

So that means if you're in the middle of one of those plateaus there's no point in arguing with this.

Because it's not going to do you any good.

We have these boundaries where we think performance break points are.

So you say, well that seems a little harsh.

Blah, blah, blah, blah, blah, and start arguing.

But then we will come back with a second major innovation we have in the course.

That is that your grade is calculated in several parts.

Part one is the max of your grade on Q1, and part one of the final.

So in other words, you get two shots at everything.

So if you have complete glorious undeniable horrible F on the first quiz, it gets erased on the final if you do well on that part of the final.

So each quiz has a corresponding mirror on the final.

You get the max of the score you got on those two pieces.

And now you say to me, I'm an MIT student.

I have a lot of guts.

I'm only going to take the final.

It has been done.

We don't recommend it.

And the reason we don't recommend it is that we don't expect everybody to do all of the final.

So there would be a lot of time pressure if you had to do all of the final, all five parts of the final.

So we have four quizzes.

And the final has a fifth part because there's some material that we teach you after the last date on which we can give you a final by institute rules.

But that's roughly how it works.

And you can read about more of the details in the FAQ on the subject homepage.

So now we're almost done.

I just want to talk a little bit about how we're going to communicate with you in the next few days, while we're getting ourselves organized.

So, number one-- if I could ask the TAs to help me pass these out-- we need to schedule you into tutorials.

So we're going to ask you to fill out this form, and give it to us before you leave.

So you'll be hearing from us once we do the sort.

There's the issue of whether we're going to have ordinary recitation and a mega recitation this week.

So pay attention.

Otherwise, you're going to be stranded in a classroom with nothing to do.

We're not going to have any regular recitations this week.

Are we having regular recitation this week, [INAUDIBLE]?

No.

We may, and probably will, have a mega recitation this week that's devoted to a Python review.

Now we know that there are many of you who are celebrating a religious holiday on Friday, and so we will be putting a lot of resources online so you can get that review in another way.

We probably will have a Python review on Friday.

And we ask that you look at our home page for further information about that as the week progresses.

So that's all folks.

That concludes what we're going to do today.

And as soon as you give us your form, we're through.