

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation, or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

**TOMASO
POGGIO:**

This problem of intelligence, it's one of those problems that mankind has been busy with it for the last 2,000 years or so. But 50 years ago or so, that was the start of artificial intelligence. It was a conference in Dartmouth, '62 or so, with people like John McCarthy and Marvin Minsky, who coined the term artificial intelligence. And at that time, progress was made. Progress has been made, especially in the last 20 years. I'll go through it. But they relied, really, only on computer science and common sense. And in the meantime, there are all these other disciplines which have made a lot of progress, and that are very likely to play a key role in the search for answers to the problem of intelligence.

So it was obvious that we needed different expertises. Not all in computer science, but in other ones. And so, this was the people that we put together from different labs, from neuroscience, from computer science, from cognitive science, and from a number of institutions in the US. Especially MIT and Harvard. Let me tell you a bit more about the background here.

This idea of merging brain research and computer science in the quest to understand intelligence. Part of the reason for this was progress and convergence we saw between different disciplines. And one of them was progress in AI. And this started, really, with Deep Blue, I guess it was called at the time. The machine IBM that managed to beat Kasparov at chess for the world championship.

And then, of course, there was Watson beating champions in Jeopardy. And things like drones able to land on aircraft carriers. So that's the most difficult thing for the pilot to do. And in the meantime, things had continued to go pretty fast. This was the cover of Nature, probably eight months ago or so. DeepMind, which is one of our industrial partners in the center, has developed an artificial intelligence called DeepQ I think, that learned to play better than humans, 49 classical Atari games. By itself. And this was two or three months ago, a cover of a Nature supplement, on artificial intelligence and machine learning.

This is showing a system by Mobileye, this is an old video, that gives vision to cars. There is a

camera looking outside, and is able to brake and accelerate when needed.

[AUDIO OUT]

There have been, there are, and there will be a lot of significant advances in AI. I think it's a golden age for intelligent applications. You know, if people want to make a lot of money with useful things, that's the time. But this is kind of engineering. Interesting one, but engineering. And we are still very far from understanding how people can answer questions about images. This is one of the main focus in the center, really. How does your brain answer simple questions about this image? About what is there? And what is this? Who is this person? What is she doing? What is she thinking? Please tell me a story about this, what's going on?

[INAUDIBLE]

And we would like to know to have a system that does that. But also, to know how our brain does it. So that's the science part. It's not enough to pass the Turing Test. In this case, to have a system that does it. We want to have a system that does it in the same way as our brain does it. And we want to compare your model, our system, with measurements on the brain of people, or monkeys, also during the same task. So that's what we call Turing plus, plus questions.

And part of the rationale about it is, this is kind of a more philosophical discussion. I personally think that it's very difficult to have a definition of intelligence, in general. There are many different forms of intelligence. What we can ask is questions about, what is human intelligence? Because we can study that. Right. You know, it is, I don't know, the ENIAC computers in the '50, more or less intelligent than a person. You know, it can do things a person cannot do. And so on. There are certain things ants or bees do, are pretty amazing. Is this intelligence? Yeah, in a certain sense is.

So I think, in terms of a well-defined question, the real question is about human intelligence. And so that's what, from the scientific part, we are focused on. And would like to be able to answer how people do understand images. We start with vision. We are not limited, eventually, to vision. But in the first five years of the center, that's the main focus. And answer the question about images. And we want to understand how the answers are produced by our brain at the computational, psychophysical, and neural level.

It's ambitious. And I think there are probably, in terms of having all these different levels, levels of really understanding from the what, where, the neuroscience, to the behavior. We are not yet at the point in which we can answer all those kind of questions at all these different levels. But some, we are. One example is, who is there? It's essentially face recognition.

And this is an interesting problem. Because we know from work, originally in the monkeys, and then with fMRI in humans. Shown here, parts of the brains of cortex, which are involved in face recognition and face perception. And then, it's possible to identify analog regions in the monkey. And record from the different patches in the monkeys, each one probably around 100,000 neurons, maybe 200,000 or so. And look at their properties when the monkey is looking at the face. And make models of what's going on. And, of course, we want these models to respect the neural data, ideally the MRI data. And do the job of recognizing faces as well as human do.

So we are getting there. I'm not saying we have the answers, but we have at least models that can be tested at all these different levels. So that's kind of the ideal situation, from the point of view of what we want to do in the center. Now as I said that, not all problems are mature at this level. There are certain like telling a story. We don't know exactly. We cannot record yet from neurons in the monkey, when the monkey is telling a story. Because the monkey has not been able to tell its story, right. And so there are other questions that are not as advanced as this one. But other type of studies can be done on them, should be done. And this is what we'll hear about.