



FAST TALK

Ayelett Shani / Photos by Gali Eytan

CAN A MACHINE BE HUMAN?

Talking with: Prof. Gal Kaminka, 42, married and a father of three, researcher in robotics at Bar-Ilan University's Computer Science Department and Gonda Multidisciplinary Brain Research Center; heads the university's MAVERICK group (Multiple Agents, Virtual Environments, Robot Intelligence, Cooperation and Knowledge) **Where:** In his lab **When:** Thursday, 2 P.M.

How much has the field of artificial intelligence progressed in the past decade?

We've achieved incredible successes with artificial intelligence (AI) that impact our lives every day. A robot that beats the world chess champion – today that seems like the most trivial thing, but people used to think that chess was the most challenging thing and that you couldn't go further than that.

Really? Intuitively, it seems to me that chess is relatively easy to decipher compared to, say, the processes of group decision-making.

True, but that's retrospective wisdom to say that chess is easier than soccer, for instance. Soccer in robotics is very hard. But in 1956, when they started with it, they didn't think so. Often, it may seem that the fathers of the field were arrogant in the declarations they made. I think they just didn't grasp the magnitude of the problem they were facing.

You entered the university today, the system identified your number, compared it to a number it had recorded in the system, and the gate opened. I don't think any of the people who built this system realized that it does something that once was considered straight out of science fiction. Computerized optics is a field that began within AI, became a scientific field unto itself, and today has made such great achievements that it is built into commercial systems, too.

What do you mean when you say the fathers of the field didn't grasp the magnitude of the problem they were facing. What's the problem?

The more we progress, the more we find that the narrower and more defined the problem – like chess, for instance – the more successful we are. We can see what works, what doesn't work, and advance. We can't yet manage to do something broader: to create something that can play chess by computer and then get in a car and drive, and then open the door too.

This sort of thing used to be thought of in terms of the Turing Test: Can an AI software program fool someone into thinking that they're dealing with an actual person?

When you really think about it, the Turing Test is an invalid one. But still it raises some interesting philosophical

questions. Are we capable of creating behavior that an outside person, a human being, will judge to be human? And we've passed this test many times within a specific context. We've built simulators in which the simulated pilots performed so well that the military personnel thought, at least for a while, that there was an actual human pilot there.

This can be done only when the field is kept very narrow. So how would you define this missing factor?

I don't know. I think we're missing something in the architecture. How the brain works, what its parts are, and how they communicate with one another in order to obtain this product that we call intelligence.

Can you give me some examples of applied research projects that you're working on – with the Defense Ministry, for example?

In recent years we've been working a lot on robotic patrols – how a team should behave on patrol, how it should move along a fence in order to identify infiltrations, and so on. Another example: In the army they use small robots that broadcast video. These can be put into the field ahead of the force to find explosives or hostile elements. In collaboration with a company called Cogniteam, we built a combined software and hardware system that can be put on such a robot, and then a single operator can control several robots at once and obtain a precise architectural sketch of the site. The robots work together as a team; they divide up the tasks and the areas, they coordinate and respond to orders, but they have the autonomy to choose where to go, unless they are told otherwise. When several robots are working in tandem, the mapping takes much less time and we lessen the risk for the soldiers who are waiting outside.

How did you get into robotics?

I was, and still am, a big fan of science fiction and [Isaac] Asimov. Not long ago I realized a dream and bought this head that you see here, of one of the first robots ever. As a kid I thought robots were really cool, and as an adult I did my doctorate in AI, and through that I came back to robotics. What interests me is intelligence in general – in animals, in people and in robots. Today, I deal with

the parts that comprise the social brain.

Robots can learn to function in society?

That's what I want them to do.

Will we see things from the world of social psychology, like social proof, where robots are concerned?

That interests me very much and doesn't exist at all in robots. On the applied level it could happen. There's a company called Kiva Systems that builds robots and was sold to Amazon for hundreds of millions of dollars [\$775 million, in March 2012]. Think about it – at Amazon you've got a vast range of products, each order is totally different, and before the robots came along the workers had to walk miles in the warehouses to collect things for each order. Kiva create a system of robots that go to the shelves and bring the products to the workers. Hundreds of robots are being operated inside a warehouse this way and they don't collide. So you think, Wow, there's social intelligence here.

But I don't think it's sufficient. If they were human workers, would you have to explain to them what to do so as not to collide? We have intuitive understanding, and what interests me is getting at these mechanisms that enable brains to work together with other brains.

And how will this happen?

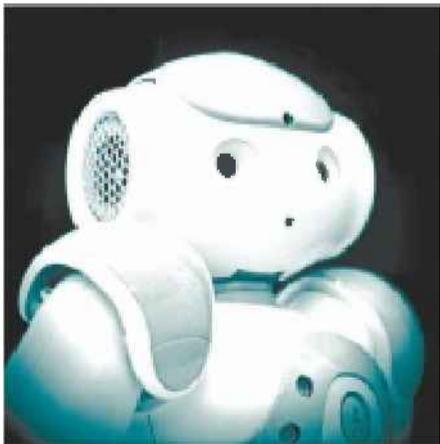
I don't know. I wrote an article about the idea of looking at studies done in social psychology and neuroscience to try to understand what differentiates people with autism from people without autism. And I thought this could be looked at as a metaphor. We program robots to behave correctly within very narrow contexts, and it's very similar to how you teach a person with Asperger's syndrome, for example, to behave in a certain social context.

That's a spot-on comparison – because they are taught these abilities without any expectation that the behavior will be internally motivated.

Exactly. I don't know how to translate it yet because the levels of abstraction are so different. But I really believe that such mechanisms of social connectivity exist, and robots just don't have that. It doesn't come from inside.

What about autonomous thinking?

We want to make autonomous thinking



26.08x29.11	2/3	עמוד 25	haaretz - haaretz magazine	04/09/2013	38639773-6
80039					



'What interests me is intelligence in general – in animals, in people and in robots. Today, I deal with the parts that comprise the social brain.'

possible, but we don't want the robot to be totally autonomous. Do you want your vacuum cleaner to decide that it doesn't feel like vacuuming today? But wait: Maybe the vacuum decides it doesn't feel like working because it consulted with the computer from the electric company and decided it would be cheaper to work at night? How many people could deal with this? Very few, I think.

How do you understand the concept of autonomous thinking in robots? I think most people have this notion that it would be connected somehow with Judgment Day or the day when robots take over the world.

We're witnessing autonomous thinking by machines all the time. Why isn't anyone upset that the washing machine is autonomous? It performs decision-making processes, such as how long to continue the cycle depending on the amount of dirt. That's decision making any way you look at it; it just doesn't

seem critical to us.

Why does the idea of robots having autonomous thinking disturb us? Is it the Frankenstein syndrome?

Possibly. Once I tried to argue that it's not scary. Then I came to the realization that there's no point fighting people who are genuinely fearful of this.

Are you fearful?

No. I was asked once what would happen if robots take over. I asked what the questioner meant by that, and he said, "What if we become dependent on the robot to the point where we can't manage without it?" I answered: "Could you turn off your telephone now for a week?" He said yes, but it would be very hard. So does that mean the telephone has taken control of you? If your Windows crashes, hasn't it gone out of control? It's annoying, but we're not claiming it tried to take control over you. That's humanization of the robot.

True, and it's very hard to resist the

temptation to do so. Why is that?

None of the robots on the market today look like a human. So why do people take the Roomba [a cleaning robot] with them on a family trip as if it were a pet? I can't understand it.

I guess that for you it's the opposite. Because you work in this field and know it so well, you're not susceptible to the rather naive and childish perception many people have of robots, a perception that is fed by popular culture.

I don't want to judge this perception, which comes out of a certain cultural context – from the Golem of Prague to Frankenstein. We grew up on tales of paradise and of technology that is bigger than us. It's a Judeo-Christian concept, a narrative that exists in our culture, in which there are things that are forbidden and shouldn't be known or touched, otherwise they'll take over or take revenge on us – all the clichés about artificial intelligence that you find in science-fiction

movies. But when you look at ordinary day-to-day life, that's not what's happening. The question is how people function when a new technology comes along, and when I look at history, I'm optimistic.

Why? We've invented lots of ways to kill ourselves with technology.

Such as?

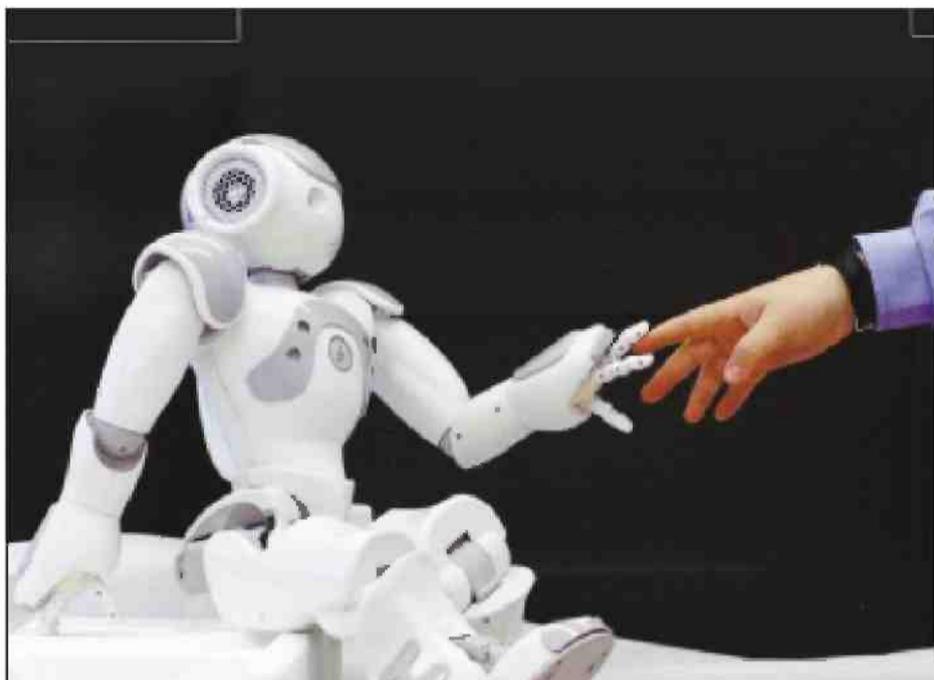
Cars. Cigarettes. Tanning beds.

The car isn't killing us. We're giving ourselves more power to hurt ourselves and others. In the end we see, throughout history, that technologies developed, and then people found creative as well as terrible ways to use them.

In what creative and terrible way can I use a robot?

In exactly the same creative and terrible way you can make use of your washing machine. I think that when talking about the ethics of robots, one has to think about the ethics of the use of machines. Someone once asked me what would happen if his wife fell in love ►

13.13x31.93	3/3	26 עמוד	haaretz - haaretz magazine	04/09/2013	38639778-1
80039					



◀ with a robot. How am I supposed to answer a question like that? Someone else once said that I treat robots as machines and it reminds him of parents who suppress their child's desire to grow and develop.

You see! We have a dialectical attitude toward robots. On the one hand we want them to be similar to us and have human qualities, and on the other we're afraid of them. And that's before we've even begun to talk about singular robots or cyborgs.

I think what you keep trying to ask, in different ways, is what differentiates the human from the machine.

Yes. And what's the answer?

That we don't know. And we won't know for a long time to come. You don't know what makes another person human. We call it the soul and the mind, but we don't know. We can say that a human is something that's not covered in fur and is a little similar to an ape, but does that capture the essence of a human being? Therefore, when I'm asked these questions, I give a simplistic answer. Can robots kill me? No. Can people use robots to kill me? Yes. Can people use the block of paper you're now holding in your hand to kill me? Yes.

We're approaching that moment. So let's return to the original question: Can a machine be human?

The answer is yes. And we have proof – us. We are a machine that is human.

That's sophistry.

It's not sophistry at all. It's a very strong statement because it rules out the divine spark, or the nonmaterial body within us. It's a very disappointing answer for me. Really. I'd like to think that there are things beyond the material.

Did you always think that way, or was this an insight you came to because of your work in artificial intelligence?

I'm still pondering this question. I'll give you the answer as I understand it now. There were periods in which, when I thought about intelligence, I formulated it to myself in terms of finding the answer. Look at the philosophy of Descartes: He talks about the separation between body and mind, and this is a way of thinking that's especially suited for people who come from the pure computer sciences, where what interests them is really just the software. The intelligence without the body.

The more I get into robotics, the more I see the influences of the body on the way of thinking. When you're irritated or scared or excited, your way of thinking changes. When you're in a panic

you run for the exit you're familiar with, even if there's another one right in front of you. In computer science terms, someone took the processor and changed it from a Pentium 3 to a Pentium 1 while your operating system was running. Emotions have a deep connection to our physiology. Stress is a very physiological thing.

Or falling in love. In each case we're unable to distinguish between the physical arousal and the emotions and thoughts that it spawns.

Exactly. And it's very hard for me to argue, given what I see in people and in robots, that thought is really something that's detached from everything else. And when you start to think about thought as something that's not detached, it's hard to make the separation and say that emotions are not detached from the body but the pure soul is.

Thought is not detached, that's clear. It's also why robots can make decisions much better than we can.

True. In narrow areas they can already make much better decisions than we can. In many cases, automatic pilots will respond in a much more efficient, quick and correct way than human pilots. I'm 42 now. I have 25 years to go until retirement. I can't tell you today what makes us unique as human beings. And the more progress I make, the more I suspect that it's not something lofty and spiritual, and not physical either.

So it's...?

Something that's related to your machine, that we don't yet understand how it's created, but that doesn't mean it's not mechanical. It's not from above and not from the hidden world. It comes from inside of us, from inside this machine that's called a human being.

I once had a literature teacher who analyzed poems in a very mechanical way. For a while it really bothered me, and then I learned to make the mental switch between just wanting to enjoy the poetry and wanting a technical understanding of how all this beauty is created. But both of these states of being are right. When you look at a picture of a beautiful rose you admire the beauty, even though you could also go over it pixel by pixel and change what shade it is.

And these – very moving – words are coming from someone who's trying to build a brain.

Who's been trying for more than 15 years to understand how to build a brain. Basically, I build robot brains and keep trying to understand. ■