

HAARETZ magazine



Violence is golden?
How radical should animal activism be? **12**

All the world's a cage
Staging Shakespeare in an Italian prison **22**

Ground forces
The best kubbeh in Israel **29**

FORTUNE TELLER

Dr. Kira Radinsky is developing a tool for predicting the future.

Next on the to-do list: saving the world

Ayelett Shani 16



FAST TALK

Ayelett Shani / Photo by Tomer Appelbaum

LIFE BY ALGORITHM

Talking to: Dr. Kira Radinsky, 27, a computer scientist specializing in artificial intelligence who develops AI systems for predicting the future. Married, lives in Zichron Yaakov **Known for:** Being on the MIT list of 35 innovators under 35 and winning awards from Google, Yahoo and Facebook **When:** Tuesday, 6:30 P.M. **Where:** Offices of the startup SalesPredict

How can the Internet be used to predict the future?

The amount of information existing in the world today is constantly expanding. Even though most people believe that this kind of information – such as Facebook searches – is rubbish, to me it is the most exciting possible thing. I fantasize about what could potentially be done with all this data. That is my passion. What I am trying to do is take all this data and construct models that attempt to predict future events.

And it worked. You've notched some impressive successes in prediction.

I always cite the Mark Twain adage, "History doesn't repeat itself, but it does rhyme." So I say that in every event, even if it possesses characteristics that are unique to it, there exists some sort of pattern drawn from the past.

Give me some examples of predictions that came true.

The riots in Sudan. Right now I am preparing a PowerPoint presentation for the prize-award ceremony at MIT, and I thought of using this as an example of a prediction. You can see a recurrent pattern in the riots: The price of a staple begins to climb – in Egypt it was bread, in Sudan it was cooking gas – and this sparks riots. The system predicted that the government in Sudan would be weakened and could even collapse.

We predicted the cholera outbreak in Cuba. We predicted the riots in Turkey, Syria and Egypt. The increase in the price of wheat. The outcome of the trial of a big-time drug dealer. I have a list of thousands of examples.

Your detractors would say it is a matter of luck. To what extent do you think luck plays a part in these predictions?

During the Second Lebanon War, my mother – a mathematician – tried to predict whether a missile would fall on her building. She joked that a missile had a 50 percent chance of hitting her building. Either it would hit it or it wouldn't. I believe that, in general, there is an element of luck in life. In the end, luck also factors into probability. The probability always exists that you will get lucky.

These predictions are based on inconceivable amounts of data. I read one estimate that there are one billion messages posted on social networks every 48 hours.

Yes. We collected 150 years' worth of news items. We took The New York Times. You can surely imagine the complexity of taking text and conducting an intensive semantic analysis of it, as if a person were reading it. We collected billions of Twitter tweets, millions of Web

searches. On top of all this data, we began to add more data that people consider obvious: encyclopedias, information on the genome, all the human data that has accumulated and that is more or less static, and then we started to make generalizations.

What do you mean by generalizations?

Let me give you a real example, one that helped me buy an iPod at a low price. Following the earthquake in Japan, the system warned of an increase in the prices of Apple products. It turns out that Apple's factories were situated in the region of the disaster and that one of them, which supplies chips, was unable to supply them. A shortage resulted, and the prices went up. The point is that we have a very large quantity of data, and we don't have to understand it deeply. Only to reach a level of generalization. That is what a human being does. He reads things and then generalizes.

Basically, the object is to imitate processes of bottom-up and top-down in human perception.

Exactly. Now, what's the problem? Let's say a person read all of this data – would he remember it? No way. Let's say, it is known that cholera erupts after storms. It isn't surprising, because this is a disease borne by water. But the system went through all of the potential combinations and managed to find a connection between outbreak of the epidemic and the famine that existed 18 months to two years before. An individual has no chance of reaching such a conclusion, because you would have to go through all the previous events and then find causal connections.

So that is a good example of the advantage a system has over human intelligence. What is the disadvantage?

The disadvantage, as usual, is that sometimes the system makes the sort of generalizations that, when I see them, I say, "Give me a break." The system often offers a prediction that is super-trivial. That after an earthquake there will be a tsunami, say. Tell me something I don't know.

And how do you get around that?

The phenomenon of two words that frequently appear one after the other is called textual entailment. The idea is that if someone knows about A, he will assume that B is also true. You read that an earthquake happened; you will assume there will be a tsunami, too.

We try to find things that did not frequently occur contiguously and that are unknown, and for which a link between them is not written about anywhere. So

the idea is to attempt to find something that filters all the textual entailment.

You work with a lot of organizations, including the Gates Foundation. What exactly are you doing with them?

Our relationship with the Gates Foundation is still fresh. Right now we are at the pilot stage, in the context of prediction of diseases and catastrophes. We are also working with the UN on the issue of preventing genocide.

Are you working with intelligence organizations?

I am not willing to respond on the topic of intelligence organizations.

What are your conclusions regarding historical determinism? Or determinism in general?

We are basically trying to spot patterns that tend to occur, to understand their significance and use them to understand which events will develop from



one another. Some of these events are related to the natural order of things, so they are determined, and some are not. For example, the system observed that riots tend to break out after policemen have shot and killed a member of a minority, and after a senior political figure – usually the mayor – has justified the police action. We saw this happen in 1991, 1999, 2000 and 2001. But in the most recent incident – the murder of the teenager Trayvon Martin in February 2012 – President Obama did not defend the actions of the legal authorities [Martin was shot by a Neighborhood Watch patroller]. He spoke of “a nation of laws.” And in my opinion, that was very beautiful. I think it may be this that prevented riots.

I wonder to what degree the desire or need to predict the future has to do with a need for control.

I think that a substantial number of critical decisions are not based on real information, simply because we have not had the means to do it. Now we can supply decision makers in the political or business realm with tools to reach better decisions.

And at the personal level? How can this technology change our lives?

Think of “my future self.” I am now at a certain stage in my career and my life. These are the things that interest me. These are the things to which I aspire. Now, let’s say that it were possible to build a system that goes through people’s Facebook pages. I can certainly find someone similar, but who is 10 years older than me, who has gone through the same things. I don’t even have to really speak with him or meet with him, only to see the course of his life and career. Or an entirely different use – let’s

assume there is someone who is 16 years old, whose friends on Facebook suddenly stop talking to him, and he is writing sad messages. Isn’t that an indicator for a potential suicide?

How far is this going to go?

My dream is that it will be possible to reach a personal prediction. A sort of brain-machine interface that goes with you, and like a GPS, tells you, “You should go to eat right now at such-and-such restaurant, because a certain individual is there now, and you will have an interesting encounter.” Just think how it would be possible to introduce this data into your personal life – that all the things you do would be based upon data.

A little frightening. Utilitarian. That is a world in which there is no form of coincidence or surprise.

We could add randomness to the algorithm.

Hilarious! Don’t you like randomness, Kira?

In fact, I do. I love surprises and enjoy them a great deal. I don’t like not knowing. That is something else. If there is something that I can know but I don’t, it annoys me, because I could have figured it out, but I didn’t have sufficient data or ability to calculate. I don’t have adequate knowledge of politics or medicine or history. I haven’t seen all the events in the world; there is simply too much data. And my life’s goal is to build the tools to cause this data to be readable. My greatest dream is that some of the data be accessible to me on a solid-state drive. And at every moment, instead of me having to do a Google search, it would be available to me.

Why not search on Google?

Because I want immediate access to this data. Most data comes to us by ►

◀ “pull”: you ask for something specific and receive it. But what if you don’t know what to ask? I think we cannot even imagine what we are capable of pulling. People don’t even know to search, and that is also a problem. The world is beginning to reach a state of affairs in which there is so much information, we no longer have the patience to even look at the search results. We want it faster. I aspire to a situation in which I don’t even have to ask the question.

Meaning what?

When you want to remember something that once happened to you, do you ask yourself, “What happened to me on such-and-such date?” No. It is immediately pulled up. And I want to know how we can get to the point where both the data and the predictions will be directly pulled up.

From the machine.

Yes, but I look at the machine as my extension. Someone once told me he didn’t want to become a Cyborg. I told him: “You’re already a Cyborg.”

The truth is, you’re right. It’s already happened. We’ve lost a lot of our abilities, and we are dependent on smartphones.

Correct. Does it make a difference that the iPhone or the iPad are not connected directly to you? When all is said and done, we no longer really function without them. I want to reach the point where it is much more intuitive.

That’s a lovely vision. Tell me about yourself, about the home you grew up in.

I immigrated to Israel from the Soviet Union when I was 4. I grew up in Neshar, with only women in the house; my mother, my aunt and my grandmother. I attended regular schools, but my mother worked three jobs so that I could receive more enrichment. Another after-school science class. When I began my studies at the Technion [Israel Institute of Technology, Haifa], I was 15; she was constantly driving me around, between her own jobs.

She must be very proud of you now.

My mother is very proud of me. For sure.

And your father?

I am not in contact with my father.

Does he know anything about you?

About your success?

I don’t know. They divorced before I was born.

So at home you were taught to excel.

Studying and education were very important to my mother. She is a mathematics expert. My aunt dealt with computers. If I was interested in something, all of them would volunteer to help. At some point they could no longer help me with the things I was interested in, and there was a need to find the proper frameworks for me.

For instance, I was very interested in biology. Particularly, genetic engineering. My mother found some summer program abroad, and somehow scraped together the money for it. I travelled there and learned how to conduct research, learned what a scientific trial looks like, how you process the results. It was very interesting to me.

You also have a black belt in karate. When did all this happen? Do you feel a need to be the best in karate, too?

I started with karate when I was 5. I simply enjoyed it. I really try hard to finish things and to be the best I could be in the things I like to do. So the black belt was a given. Part of my interest in karate is discipline. I like it. I learned a lot about



Riots in Tahrir Square, Egypt, January 2011. A recurrent pattern involving the price of a staple prefigures many riots.

patience and honor, and doing the best you can.

Let’s talk about the privacy issue.

Everyone is so frightened about this matter of privacy. I really think that we have much more to gain than lose. Say I want to share my genetic data. I want to give it to scientific research. That option does not exist. Look at companies like [U.S. genetics-testing firm] 23andMe.

You are referring to a technology that enables the prediction of diseases – you give the company a sample of your saliva, they analyze the DNA, and they predict what diseases you might develop.

Yes. And what is so cool about this concept is that they compare your findings with a very large human database.

The trend toward sharing of medical data is growing stronger. But who says it is really for the best?

I don’t understand what people are afraid of. Really. It is possible to enact regulations for this thing. There was a time when people used to think the brain of a person traveling on a train at a certain speed would explode.

It really isn’t on the same level. It is well known that corporations exploit the data for negative purposes, and I don’t think that regulation would help. In any case, most of these things are done clandestinely.

It’s already there. Let’s see: How much data is there about you on Facebook?

I’m not active on Facebook.

Then for the majority of people, all of the data, the photos and the habits, are already there. Everything is tagged. Facebook knows exactly with which friend I spent time. This is, incidentally, one of my favorite predictions. Once, I tried to determine the likelihood of me coming down with swine flu. The outbreak began in Mexico, so I looked for people who’d travelled there. Who was in the pictures with them? Did they meet with anyone with whom I had met? Within a few steps, I could pretty effectively predict how many people on my team would get sick and would not be coming in to work. It was a cute project.

‘I don’t like **not knowing**. If there is something that I can know but I don’t, it annoys me, because I could have figured it out, but I didn’t have sufficient data or ability to calculate.’

So I would say that no, there is no longer any privacy.

Correct. Privacy is an illusion, and the corporations can make negative use of this data.

Okay, so why don’t they do it already?

Clearly, they are doing it. Do you know what they do with all of your details and data? Even the most innocent uses are pretty shocking. If I search for a flight, is it to my benefit that for the entire week afterward I am being offered advertisements about flights?

Correct. And if they offer you a discounted flight, is that something bad?

No. But if I have a gambling problem, and they are incessantly urging me to come and play poker, then it is serious.

So I am saying the opposite. Do you have a gambling problem? This is an opportunity to resolve it. I had the occasion to speak with [President] Shimon Peres about precisely these matters. He said something very wise: that he sees that governments are being brought down from their places on high, and the new form of government is commercial companies. But large corporations are not all evil. Let’s say that we want to prevent crime. Many types of crime start with gambling. So someone can see that someone’s gambling problem is developing, based on the fact that he has a large

number of searches related to gambling. So send him all sorts of things that would help him resolve the problem.

Sure. Google will cure me of my gambling problem.

I think that if I would go to Google and say to them, “We want to change the world, we want to prevent crime all around the world, and you have all of the data that can help us do so,” they would not be able to do a thing, for the reason that they are restricted by privacy issues. It would help if the restrictions on them were rescinded. I know of a research study conducted at Microsoft not long ago. They took people’s searches and attempted to identify psychotic attacks. It emerged that when people who are known to suffer from mental illness forget to take their medication, their searches change. For instance, they begin to search for all sorts of things of a sexual nature. It is at that point the physician could be informed that things are beginning to unravel.

These are truly complex issues, of practical ethics.

I think that a better society can be built by means of this data. It would not of necessity be bad.

I am compelled to represent the bad side. Negative use could also be made of your technology. First of all, who actually has access to the technology?

It is all on open source. All of the data, everything I wrote in my doctoral thesis – it’s all out in the open.

If I am a large, evil food conglomerate, and I know that a shortage of rice can be expected in the Philippines next year, I could potentially conduct a lot of manipulations on the basis of this data.

Correct. And that is what is happening, say, in algo trading [automatic trading of shares, carried out by algorithm]. They start to compete against one another and cause all sorts of problems. One algorithm will suddenly knock out all the others, or a dependency between algorithms is formed. But that is a new type of economy - what’s wrong with it? It is merely a different economy.

The point is that there is no way to predict the consequences.

Is nuclear energy a good thing or a bad thing? It depends on how it is used. With all technology, you can do something good and something bad.

But what happens when the technology meddles with reality? Let's say you predicted a certain event, and measures were taken to prevent it. The technology is essentially creating an absolutely different series of events.

The system is not static. It learns from what is happening and places less weight on things that happened in the past. If the pattern it found is then broken, it begins to find other patterns. The system adjusts itself.

Does it learn how to weigh the prediction effect into itself?

It does not know this is an effect of the prediction; that it affected the reality. It simply sees that there is no drought. You prevented it? Great. You will be healthy.

How did you even think of this? Why did you decide to try to predict the future?

It was a process. Google publicized all the questions people feed into the aggregator, and you can really see how it behaves over time. I told my husband, "Look how cool this is. It has a correlation with events as they are happening, and there are all sorts of methods of predicting the next point on the time series. How cool this is – you can predict how it is going to develop." And he said, "Cool. So why don't you try to predict it with the help of other events and correlations?" And I said, "Yes, it would be pretty simple to do that." I tried, and it really did work.

How old were you at the time?

Twenty-one. I started to do it and then I said, Hold on, why should I only use the questions posed to Google? I have so many different forms of data; there is news from around the world. And then I got into the project on a grander scale.

And this, in fact, is your doctoral thesis. Some people thought you were megalomaniac.

Yes. I worked on it with my adviser, Prof. Shaul Markovitch. It was very difficult to prove that my examination was good, not that the algorithms were good. It's always very difficult to publish an academic paper. There is a lack of understanding, there are ego wars, but as soon as you can persuade someone that it works, then it flows freely.

How did you persuade them?

I submitted it several times, each time correcting what I had to. In the end, it is on the level of "These are the events we predicted, and this is what happened." It is pretty hard to argue with that. The first time, I was told, "Yes, but even if you hadn't predicted it, it is trivial that it did occur." So I said, "Okay, here are X number of events, you say what you think will happen, and the system will say what will happen, and let's see who is more right." And then we saw that the system was predicting significantly better than people.

Explain a bit about how learning algorithms work.

Let's say we take a group of people and we know what their attributes are – height, eye color, address, etc. For each of them, there is a notation of whether they are rich or poor. Should a new individual arrive, about whom you do not know if he is rich or poor, but we want to know if he is rich or poor, the system is a black box. It does something inside, and

when someone new shows up, it looks at the parameters, determines if he is similar to a rich or a poor man, and produces a probability. That is how the black box works, and entire bodies of research are devoted to constructing these black boxes.

In what way is the technology you developed different from what other companies are doing? For instance, Recorded Future, which, among other things, predicted the crisis in Yemen in 2010.

Recorded Future gathers data on what people think will happen in the future. For example, they gather rumors from the Internet regarding the next iPad launch, and make an aggregation of it all in order to produce a prediction. I think their technology is very interesting because they are using the wisdom of the masses to produce the prediction. The technology we've developed attempts to find patterns that have not previously been observed in digital history, and it produces new probabilities.

Can you use this technology not only to predict the future, but also to draw conclusions on how crises are managed in the present? Would it have been possible, for instance, to deduce from the shutdown of federal services in the U.S. in 1995, during the Clinton administration, how to manage the recent shutdown?

I think so. It is possible. And not only from one shutdown of an American administration to another, but also from similar shutdowns worldwide from which one can learn. The system we developed is capable of avoiding oversimplification, and finding only the recurrent pattern. For example, it will not predict after the shutdown that the Queen will dissolve Parliament, as happened in Australia. [In 1975, following a similar shutdown, the Governor-General who was considered the delegate of Queen Elizabeth, Sir John Kerr, dismissed the Australian prime minister and members of parliament.] Due to the fact that the system knows there is a substantive difference between the U.S. and Australia, at least in this sense the U.S. is not subject to the Queen's authority.

You founded a startup that utilizes this technology for commercial purposes.

It isn't exactly the same technology I developed in my doctorate; that was based on concepts. Here we are basically helping companies to improve their sales.

How does it work?

Let's say that a company like HP, which sells servers, approaches me. They do not want to waste more resources on approaching customers who don't really want to buy. Moreover, they want to know how many servers they need to keep in inventory. At the end of the day, every CEO is measured according to this: To what degree he succeeded in predicting, as opposed to what in fact happened. If the stock falls, it is due to these things.

So what do you do?

I look for examples of companies that purchased servers in the past. And examples of companies that were not interested in purchasing, and all of the data HP has on these companies. Transcripts of conversations, for example. In addition, the company gives me a list of companies that it has not yet approached, and like in a search engine that grades the results, I begin to rank the people. For every individual, I go to the Web and begin to find out how senior he is in the firm; how large the firm is; the degree to

'I immigrated to Israel from the Soviet Union when I was 4. I grew up in Neshet, with only women in the house: my mother, my aunt and my grandmother.'

which this firm is becoming more professional. We carry out a full-blown analysis of the firm's website, where its business is coming from, how much money each of its divisions is earning.

Does the CEO's wife have her own server factory?

That is a little more complicated to find, but you can see, say, if he is tracking his competitor on Twitter.

Again, most of our work today is creating parameters. And the system finds links that are hard to explain. Let's say we found that there is a greater chance of an insurance company purchasing servers if the data shows that surfers are remaining on this insurance company's website for a long time.

Why is that interesting? Think of what an insurance company's website looks like. If people stay on the site for a long time, they are evidently buying insurance, so if this insurance company has a lot of customers, then it has a lot of money for technology and it will invest in servers.

A lot of these things are difficult to explain to people. It's like with the drought. I call it hidden variables.

What sort of information do you tap into in order to conduct these analyses, aside from the technological data? Consumer behavior? Behavioral economics?

For every company – even for every product – these are different associations. At the end of the day, these algorithms make up a lot of their own rules.

Because they learn from themselves.

Yes. They can amount to thousands of laws. I cannot go through thousands of laws. It reaches levels of micro-segmentation that I can't really comprehend.

Predictive analysis has become very trendy, with lots of firms trying to do it.

That makes me happy. I truly believe that we have finally reached the critical mass of data and calculating power that are required to start predicting the future systematically.

On the philosophical level, as you see things now, what does human intelligence have that artificial intelligence cannot reproduce? What does human intelligence have that a machine would be hard-pressed to replicate?

It will sound a little funny to you, but it's the trivial things. Let's say I recite a line like "I threw a glass against the wall and it broke." You instantly understand that the glass broke, but the computer doesn't understand that. It does not know if I meant the wall or the glass. And if I would give you additional data and tell you that the wall was glass and the cup was steel, then you would immediately understand that the wall broke. The computer wouldn't know. Things that a 2- or 3-year-old child already knows are very difficult for the computer. And that is where most of the work is. All the

things we learn from interaction with the world – the computer doesn't know them. Think of something that never existed on earth, and you have to explain to it all the things that seem intuitive to you. That is so hard. Nearly impossible.

Where do you think you will be when you are 40?

I am thinking of a mix between Bill Gates and Elon Musk. Musk is one of the most incredible engineers in the world. You ask him, "Could you build a spaceship for \$2 million?" and he says, "Yes." I love his vision and the fact that he makes things happen. With Bill Gates, I love the fact that he is such a successful businessman, and much of what he has done has been done to improve the state of the world. So at age 40, I want to do things with vision, great things, but not to do them simply for money or to make a better algo trade.

What motivates you?

I have a lot of motivations, I suppose. I am interested in doing something really big. I don't feel I have done anything great until now. I want to have an effect on people's lives, in a good way.

Why do you feel you have to do something great?

I feel that is my calling.

You want to change the world.

That is my email signature. "Change the world Kira." And the first time I sent someone an email with that signature, he responded to me with something along the lines of, "Oh, c'mon!" But someone else answered, "Kira, please do." I don't think everyone should aspire to it, but it gives me a feeling that this is not just life, but life that has meaning. I wouldn't want to look backward at my life and realize I haven't done a thing.

Do your drive and ability isolate you a bit from the world?

No. I have my set of very close friends. And an amazing life partner. I have known my husband since I was 8. He is a genius. There isn't anyone else that I love to talk with more than him. I have always been able to share everything with him. And he has also supported me. In the course of my studies I was given the option of going to Microsoft Research in the United States, and he left his job, and we went there. To this day, I admire him for that.

This isn't just any job he left. Your husband, Sagie Davidovich, is also a startup entrepreneur.

Yes. His first startup was sold to Sears, and since then he's been involved in two other startups. And now he has founded a new startup. He also grew up in a home where his father nurtured his fields of interest. He sold his first software when he was in fourth grade, and bought a bicycle with the money. I remember him coming over and telling me excitedly, "Kira, I sold software!" We began working together. We were in an after-school electronics club and we invented a flashlight that switched colors.

Do you think you can achieve everything you want?

In the event that I have invested adequate time and effort, yes.

You're not embarrassed or scared? There is nothing that stops you?

I think that shame and fear deter a lot of people. I don't think about it in those terms. As far as I'm concerned, there is an objective and there are very logical ways of reaching it; I am simply planning and thinking what the best and most effective way of reaching it is. ■