

“ *The current machine learning concentration on deep learning and its non-transparent structures is a hang-up. They need to liberate themselves from this data-centric philosophy.* ”

JUDEA PEARL

PROFESSOR OF COMPUTER SCIENCE AND STATISTICS,
UNIVERSITY OF CALIFORNIA, LOS ANGELES
DIRECTOR OF THE COGNITIVE SYSTEMS LABORATORY

Judea Pearl is known internationally for his contributions to artificial intelligence, human reasoning, and philosophy of science. He is particularly well known in the AI field for his work on probabilistic (or Bayesian) techniques and causality. He is the author of more than 450 scientific papers and three landmark books: Heuristics (1984), Probabilistic Reasoning (1988), and Causality (2000; 2009). His 2018 book, The Book of Why, makes his work on causation accessible to a general audience. In 2011, Judea received the Turing Award, which is the highest honor in the field of computer science and is often compared to the Nobel Prize.

ARCHITECTS OF INTELLIGENCE

MARTIN FORD: You've had a long and decorated career. What path led you to get started in computer science and artificial intelligence?

JUDEA PEARL: I was born in Israel in 1936, in a town named Bnei Brak. I attribute a lot of my curiosity to my childhood and to my upbringing, both as part of Israeli society and as a lucky member of a generation that received a unique and inspiring education. My high-school and college teachers were top-notch scientists who had come from Germany in the 1930s, and they couldn't find a job in academia, so they taught in high schools. They knew they would never get back to academia, and they saw in us the embodiment of their academic and scientific dreams. My generation were beneficiaries of this educational experiment—growing up under the mentorship of great scientists who happened to be high-school teachers. I never excelled in school, I was not the best, or even second best, I was always third or fourth, but I always got very involved in each area taught. And we were taught in a chronological way, focusing on the inventor or scientist behind the invention or theorem. Because of this, we got the idea that science is not just a collection of facts, but a continuous human struggle with the uncertainties of nature. This added to my curiosity.

I didn't commit myself to science until I was in the army. I was a member of a Kibbutz and was about to spend my life there, but smart people told me that I would be happier if I utilized my mathematical skills. As such, they advised me to go and study electronics in Technion, the Israel Institute of Technology, which I did in 1956. I did not favor any particular specialization in college; but I enjoyed circuit synthesis and electromagnetic theory. I finished my undergraduate degree and got married in 1960. I came to the US with the idea of doing graduate work, getting my PhD, and going back.

MARTIN FORD: You mean you planned to go back to Israel?

JUDEA PEARL: Yes, my plan was to get a degree and come back to Israel. I first registered at the Brooklyn Polytechnic Institute (now part of NYU), which was one of the top schools in microwave communication at the time. However, I couldn't afford the tuition, I ended up employed at the David Sarnoff Research Center at the RCA laboratory in Princeton, New Jersey. There, I was a member of the computer memory group under Dr. Jan Rajchman, which was a hardware-oriented group. We, as well as everybody else in the country, were looking for different physical

JUDEA PEARL

mechanisms that could serve as computer memory. This was because magnetic core memories became too slow, too bulky, and you had to string them manually.

People understood that the days of core memory were numbered, and everybody—IBM, Bell Labs, and RCA Laboratories—was looking for various phenomena that could serve as a mechanism to store digital information. Superconductivity was appealing at that time because of the speed and the ease of preparing the memory, even though it required cooling to liquid helium temperature. I was investigating circulating currents in superconductors, again for use in memory, and I discovered a few interesting phenomena there. There's even a Pearl vortex named after me, which is a turbulent current that spins around in superconducting films, and gives rise to a very interesting phenomenon that defies Faraday's law. It was an exciting time, both on the technological side and on the inspirational, scientific side.

Everyone was also inspired by the potential capabilities of computers in 1961 and 1962. No one had any doubt that eventually, computers would emulate most human intellectual tasks. Everyone was looking for tricks to accomplish those tasks, even the hardware people. We were constantly looking for ways of making associative memories, dealing with perception, object recognition, the encoding of visual scenes; all the tasks that we knew are important for general AI. The management at RCA also encouraged us to come up with inventions. I remember our boss Dr. Rajchman visiting us once a week and asking if we had any new patent disclosures.

Of course, all work on superconductivity stopped with the advent of semiconductors, which, at the time, we didn't believe would take off. We didn't believe that miniaturization technology would succeed as it did. We also didn't believe they could overcome the vulnerability problem where the memory would be wiped if the battery ran out. Obviously, they did, and semiconductor technology wiped out all its competitors. At that point, I was working for a company called Electronic Memories, and the rise of semiconductors left me without a job. That was how I came to academia, where I pursued my old dreams of doing pattern recognition and image encoding.

MARTIN FORD: Did you go directly to UCLA from Electronic Memories?

JUDEA PEARL: I tried to go to the University of Southern California, but they wouldn't hire me because I was too sure of myself. I wanted to teach software,

ARCHITECTS OF INTELLIGENCE

even though I'd never programmed before, and the Dean threw me out of his office. I ended up at UCLA because they gave me a chance of doing the things that I wanted to do, and I slowly migrated into AI from pattern recognition, image encoding, and decision theory. The early days of AI were dominated by chess and other game-playing programs, and that enticed me in the beginning, because I saw there a metaphor for capturing human intuition. That was and remained my life dream, to capture human intuition on a machine.

In games, the intuition comes about in the way you evaluate the strength of a move. There was a big gap between what machines can do and what experts can do, and the challenge was to capture experts' evaluation in the machine. I ended up doing some analytical work and came up with a nice explanation of what heuristics is all about, and an automatic way of discovering heuristics, it is still in use today. I believe I was the first to show that alpha-beta search is optimal, as well other mathematical results about what makes one heuristic better than another. All of that work was compiled in my book, *Heuristics*, which came out in 1983. Then expert systems came to the scene, and people were excited about capturing different kinds of heuristics—not the heuristic of a chess master, but the intuition of highly-paid professionals, like a physician or a mineral explorer. The idea was to emulate professional performance on a computer system, either to replace or to assist the professional. I looked at expert systems as another challenge of capturing intuition.

MARTIN FORD: Just to clarify, expert systems are mostly based on rules, correct? If this is true, then do that, etc.

JUDEA PEARL: Correct, it was based on rules, and the goal was to capture the mode of operation of an expert, what makes an expert decide one way or the other while engaging in professional work.

What I did, was to replace it with a different paradigm. For example, instead of modeling a physician—the expert—we modeled the disease. You don't have to ask the expert what they do. Instead, you ask, what kind of symptoms you expect to see if you have malaria or if you have the flu; and what do you know about the disease? On the basis of this information, we built a diagnosis system that could examine a collection of symptoms and come out with the suspected disease. It also works for mineral exploration, for troubleshooting, or for any other expertise.

JUDEA PEARL

MARTIN FORD: Was this based on your work on heuristics, or are you referring now to Bayesian networks?

JUDEA PEARL: No, I left heuristics the moment my book published in 1983, and I started working on Bayesian networks and uncertainty management. There were many proposals at the time for managing uncertainties, but they didn't gel with the dictates of probability theory and decision theory, and I wanted to do it correctly and efficiently.

MARTIN FORD: Could you talk about your work on Bayesian networks? I know they are used in a lot of important applications today.

JUDEA PEARL: First, we need to understand the environment at the time. There was a tension between the scruffies and the neaties. The scruffies just wanted to build a system that works, not caring about guarantees or whether their methods comply with any theory or not. The neaties wanted to understand why it worked and make sure that they have performance guarantees of some kind.

MARTIN FORD: Just to clarify, these were nicknames for two groups of people with different attitudes.

JUDEA PEARL: Yes. We see the same tension today in the machine learning community, where some people like to get machines to do important jobs, regardless of whether they're doing it optimally or whether the system can explain itself as long as the job is being done. The neaties would like to have explainability and transparency, systems that can explain themselves and systems that have performance guarantees.

Well, at that time, the scruffies were in command, and they still are today, because they have a good conduit to funders and to industry. Industry, however, is short-sighted and requires short-term success, which creates an imbalance in research emphasis. It was the same in the Bayesian network days; the scruffies were in command. I was among the few loners who advocated doing things correctly by the rules of probability theory. The problem was that probability theory, if you adhere to it in the traditional way, would require exponential time and exponential memory, and we couldn't afford these two resources.

I was looking for a way of doing it efficiently, and I was inspired by the work of David Rumelhart, a cognitive psychologist who examined how children read text

ARCHITECTS OF INTELLIGENCE

so quickly and reliably. His proposal was to have a multi-layered system going from the pixel level to the semantic level, then the sentence level and the grammatical level, and they all shake hands and pass messages to each other. One level doesn't know what the other's doing; it's simply passing messages. Eventually, these messages converge on the correct answer when you read a word like "the car" and distinguish it from "the cat," depending on the context in the narrative.

I tried to simulate his architecture in probability theory, and I couldn't do it very well until I discovered that if you have a tree as a structure connecting the modules, then you do have this convergence property. You can propagate messages asynchronously, and eventually, the system relaxes to the correct answer. Then we went to a polytree, which is a fancier version of a tree, and eventually, in 1995, I published a paper about general Bayesian networks.

This architecture really caught us by surprise because it was very easy to program. A programmer didn't have to use a supervisor to oversee all the elements, all they had to do was to program what one variable does when it wakes up and decides to update its information. That variable then sends messages to its neighbors. The neighbors send messages to their neighbors, and so on. The system eventually relaxes to the correct answer.

The ease of programming was the feature that made Bayesian networks acceptable. It was also made acceptable by the idea that you can program the disease and not the physician—the domain, and not the professional that deals with the domain—that made the system transparent. The users of the system understood why the system provided one result or another, and they understood how to modify the system when things changed in the environment. You had the advantage of modularity, which you get when you model the way things work in nature.

It's something that we didn't realize at the time, mainly because we didn't realize the importance of modularity. When we did, I realized that it is causality that gives us this modularity, and when we lose causality, we lose modularity, and we enter into no-man's land. That means that we lose transparency, we lose reconfigurability, and other nice features that we like. By the time that I published my book on Bayesian networks in 1988, though, I already felt like an apostate because I knew already that the next step would be to model causality, and my love was already on a different endeavor.

JUDEA PEARL

MARTIN FORD: We always hear people saying that “correlation is not causation,” and so you can never get causation from the data. Bayesian networks do not offer a way to understand causation, right?

JUDEA PEARL: No, Bayesian networks could work in either mode. It depends on what you think about when you construct it.

MARTIN FORD: The Bayesian idea is that you update probabilities based on new evidence so that your estimate should get more accurate over time. That’s the basic concept that you’ve built into these networks, and you figured out a very efficient way to do that for a large number of probabilities. It’s clear that this has become a really important idea in computer science and AI because it’s used all over the place.

JUDEA PEARL: Using Bayes’ rule is an old idea; doing it efficiently was the hard part. That’s one of the things that I thought was necessary for machine learning. You can get evidence and use the Bayesian rule to update the system to improve its performance and improve the parameters. That’s all part of the Bayesian scheme of updating knowledge using evidence, it is probabilistic, not causal knowledge, so it has limitations.

MARTIN FORD: But it’s used quite frequently, for example, in voice recognition systems and all the devices that we’re familiar with. Google uses it extensively for all kinds of things.

JUDEA PEARL: People tell me that every cellphone has a Bayesian network doing error correction to minimize transmission noise. Every cellphone has a Bayesian network and belief propagation, that’s the name we gave to the message passing scheme. People also tell me that Siri has a Bayesian network in it, although Apple is too secretive about it, so I haven’t been able to verify it.

Although Bayesian updating is one of the major components in machine learning today, there has been a shift from Bayesian networks to deep learning, which is less transparent. You allow the system itself to adjust the parameters without knowing the function that connects input and output. It’s less transparent than Bayesian networks, which had the feature of modularity, and which we didn’t realize was so important. When you model the disease, you actually model the cause and effect relationship of the disease, not the expert, and you get modularity. Once

ARCHITECTS OF INTELLIGENCE

we realize that, the question begs itself: What is this ingredient that you and I call “cause and effect relationships”? Where does it reside, and how do you handle it? That was the next step for me.

MARTIN FORD: Let’s talk about causation. You published a very famous book on Bayesian networks, and it was really that paper that led to Bayesian techniques becoming so popular in computer science. But before that book was even published, you were already starting to think about moving on to focus on causation?

JUDEA PEARL: Causation was part of the intuition that gave rise to Bayesian networks, even though the formal definition of Bayesian networks is purely probabilistic. You do diagnostics, you make predictions, and you don’t deal with interventions. If you don’t need interventions, you don’t need causality—theoretically. You can do everything that a Bayesian network does with purely probabilistic terminology. However, in practice, people noticed that if you structure the network in the causal direction, things are much easier. The question was why.

Now we understand that we were craving for features of causality that we didn’t even know come from causality. These were: modularity, reconfigurability, transferability, and more. By the time I looked into causality, I had realized that the mantra “correlation does not imply causation” is much more profound than we thought. You need to have causal assumptions before you can get causal conclusions, which you cannot get from data alone. Worse yet, even if you are willing to make causal assumptions, you cannot express them.

There was no language in science in which you can express a simple sentence like “mud does not cause rain,” or “the rooster does not cause the sun to rise.” You couldn’t express it in mathematics, which means that even if you wanted to take it for granted that the rooster does not cause the sun to rise, you couldn’t write it down, you couldn’t combine it with data, and you couldn’t combine it with other sentences of this kind.

In short, even if you agree to enrich the data with causal assumptions, you couldn’t write down the assumptions. It required a whole new language. This realization was really a shock and a challenge for me because I grew up on statistics, and I believed that scientific wisdom lies in statistics. Statistics allows you to do induction, deduction, abduction, and model updating. And here I find the language of statistics

JUDEA PEARL

crippled in hopeless helplessness. As a computer scientist, I was not scared because computer scientists invent languages to fit their needs. But what is the language that should be invented, and how do we marry this language with the language of data?

Statistics speaks a different language—the language of averages, of hypothesis testing, summarizing data and visualizing it from different perspectives. All of this is the language of data, and here comes another language, the language of cause and effect. How do we marry the two so that they can interact? How do we take assumptions about cause and effect, combine them with the data that I have, and then get conclusions that tell me how nature works? That was my challenge as a computer scientist and as a part-time philosopher. This is essentially the role of a philosopher, to capture human intuition and formalize it in a way that it can be programmed on a computer. Even though philosophers don't think about the computer, if you look closely at what they are doing, they are trying to formalize things as much as they can with the language available to them. The goal is to make it more explicable and more meaningful so that computer scientists can eventually program a machine to perform cognitive functions that puzzle philosophers.

MARTIN FORD: Did you invent the technical language or the diagrams used for describing causation?

JUDEA PEARL: No, I didn't invent that. The basic idea was conceived in 1920 by a geneticist named Sewall Wright, who was the first to write down a causal diagram with arrows and nodes, like a one-way city map. He fought all his life to justify the fact that you can get things out of this diagram that statisticians could not get from regression, association, or from correlation. His methods were primitive, but they proved the point that he could get things that the statisticians could not get.

What I did was to take Sewall Wright's diagrams seriously and invested into them all my computer science background, reformed them, and exploited them to their utmost. I came up with a causal diagram as a means of encoding scientific knowledge and as a means of guiding machines in the task of figuring out cause-effect relationships in various sciences, from medicine, to education, to climate warming. These were all areas where scientists worry about what causes what, how nature transmits the information from cause to effect, what are the mechanisms involved, how do you control it, and how do you answer practical questions which involve cause-effect relationships.

ARCHITECTS OF INTELLIGENCE

This has been my life's challenge for the past 30 years. I published a book on that in 2000, with the second edition in 2009, called *Causality*. I co-authored a gentler introduction in 2015. And this year, I co-authored *The Book of Why*, which is a general audience book explaining the challenge in down-to-earth terms, so that people can understand causality even without knowing equations. Equations of course help to condense things and to focus on things, but you don't have to be a rocket scientist to read *The Book of Why*. You just have to follow the conceptual development of the basic ideas. In that book, I look at history from a causal lens perspective; I asked what conceptual breakthroughs made a difference in the way we think, rather than what experiments discovered one drug or another.

MARTIN FORD: I've been reading *The Book of Why* and I'm enjoying it. I think one of the main outcomes of your work is that causal models are now very important in the social and natural sciences. In fact, I just saw an article the other day, written by a quantum physicist who used causal models to prove something in quantum mechanics. So clearly your work has had a big impact in those areas.

JUDEA PEARL: I read that article. In fact, I put it on my next-to-read list because I couldn't quite understand the phenomena that they were so excited about.

MARTIN FORD: One of the main points I took away from *The Book of Why* is that, while natural and social scientists have really begun to use the tools of causation, you feel that the field of AI is lagging behind. You think AI researchers will have to start focusing on causation in order for the field to progress.

JUDEA PEARL: Correct. Causal modeling is not at the forefront of the current work in machine learning. Machine learning today is dominated by statisticians and the belief that you can learn everything from data. This data-centric philosophy is limited.

I call it curve fitting. It might sound derogatory, but I don't mean it in a derogatory way. I mean it in a descriptive sense that what people are doing in deep learning and neural networks is fitting very sophisticated functions to a bunch of points. These functions are very sophisticated, they have thousands of hills and valleys, they're intricate, and you cannot predict them in advance. But they're still just a matter of fitting functions to a cloud of points.

JUDEA PEARL

This philosophy has clear theoretical limitations, and I'm not talking about opinion, I'm talking about theoretical limitations. You cannot do counterfactuals, and you cannot think about actions that you've never seen before. I describe it in terms of three cognitive levels: seeing, intervening, and imagining. Imagining is the top level, and that level requires counterfactual reasoning: how would the world look like had I done things differently? For example, what would the world look like had Oswald not killed Kennedy, or had Hillary won the election? We think about those things and can communicate with those kinds of imaginary scenarios, and we are quite comfortable to engage in this "let's pretend" game.

The reason why we need this capability is to build new models of the world. Imagining a world that does not exist gives us the ability to come up with new theories, new inventions, and also to repair our old actions so as to assume responsibility, regret, and free will. All of this comes as part of our ability to generate worlds that do not exist but could exist, but still generate them widely, not wildly. We have rules for generating plausible counterfactuals that are not whimsical. They have their own inner structure, and once we understand this logic, we can build machines that imagine things, that assume responsibility for their actions, and understand ethics and compassion.

I'm not a futurist and I try not to talk about things that I don't understand, but I did some thinking, and I believe I understand how important counterfactuals are in all these cognitive tasks that people dream of which eventually will be implemented on a computer. I have a few basic sketches of how we can program free will, ethics, morality, and responsibility into machines, but these are in the realm of sketches. The basic thing is that we know today what it takes to interpret counterfactuals and understand cause and effect.

These are the mini-steps toward general AI, but there's a lot we can learn from these steps, and that's what I'm trying to get the machine learning community to understand. I want them to understand that deep learning is a mini-step toward general AI. We need to learn what we can from the way theoretical barriers were circumvented in causal reasoning, so that we can circumvent them in general AI.

MARTIN FORD: So, you're saying that deep learning is limited to analyzing data and that causation can never be derived from data alone. Since people are able to do

ARCHITECTS OF INTELLIGENCE

causal reasoning, the human mind must have some built-in machinery that allows us to create causal models. It's not just about learning from data.

JUDEA PEARL: To create is one thing, but even if somebody creates it for us, our parents, our peers, our culture, we need to have the machinery to utilize it.

MARTIN FORD: Right. It sounds like a causal diagram, or a causal model is really just a hypothesis. Two people might have different causal models, and somewhere in our brain is some kind of machinery that allows us to continuously create these causal models internally, and that's what allows us to reason based on data.

JUDEA PEARL: We need to create them, to modify them, and to perturb them when the need arises. We used to believe that malaria is caused by bad air, now we don't. Now we believe it's caused by a mosquito called Anopheles. It makes a difference because if it is bad air, I will carry a breathing mask the next time I go to the swamp; and if it's an Anopheles mosquito, I'll carry a mosquito net. These competing theories make a big difference in how we act in the world. The way that we get from one hypothesis to another was by trial and error; I call it playful manipulation.

This is how a child learns causal structure, by playful manipulation, and this is how a scientist learns causal structure—playful manipulation. But we have to have the abilities and the template to store what we learn from this playful manipulation so we can use it, test it, and change it. Without the ability to store it in a parsimonious encoding, in some template in our mind, we cannot utilize it, nor can we change it or play around with it. That is the first thing that we have to learn; we have to program computers to accommodate and manage that template.

MARTIN FORD: So, you think that some sort of built-in template or structure should be built into an AI system so it can create causal models? DeepMind uses reinforcement learning, which is based on practice or trial and error. Perhaps that would be a way of discovering causal relationships?

JUDEA PEARL: It comes into it, but reinforcement learning has limitations, too. You can only learn actions that have been seen before. You cannot extrapolate to actions that you haven't seen, like raising taxes, increasing the minimum wage, or banning cigarettes. Cigarettes have never been banned before, yet we have

JUDEA PEARL

machinery that allows us to stipulate, extrapolate, and imagine what could be the consequences of banning cigarettes.

MARTIN FORD: So, you believe that the capability to think causally is critical to achieving what you'd call strong AI or AGI, artificial general intelligence?

JUDEA PEARL: I have no doubt that it is essential. Whether it is sufficient, I'm not sure. However, causal reasoning doesn't solve every problem of general AI. It doesn't solve the object recognition problem, and it doesn't solve the language understanding problem. We basically solved the cause-effect puzzle, and we can learn a lot from these solutions so that we can help the other tasks circumvent their obstacles.

MARTIN FORD: Do you think that strong AI or AGI is feasible? Is that something you think will happen someday?

JUDEA PEARL: I have no doubt that it is feasible. But what does it mean for me to say no doubt? It means that I am strongly convinced it can be done because I haven't seen any theoretical impediment to strong AI.

MARTIN FORD: You said that way back around 1961, when you were at RCA, people were already thinking about this. What do you think of how things have progressed? Are you disappointed? What's your assessment of progress in artificial intelligence?

JUDEA PEARL: Things are progressing just fine. There were a few slowdowns, and there were a few hang-ups. The current machine learning concentration on deep learning and its non-transparent structures is such a hang-up. They need to liberate themselves from this data-centric philosophy. In general, the field has been progressing immensely, because of technology and because of the people that the field attracts. The smartest people in science.

MARTIN FORD: Most of the recent progress has been in deep learning. You seem somewhat critical of that. You've pointed out that it's like curve fitting and it's not transparent, but actually more of a black-box that just generates answers.

JUDEA PEARL: It's curve fitting, correct, it's harvesting low-hanging fruits.

MARTIN FORD: It's still done amazing things.

ARCHITECTS OF INTELLIGENCE

JUDEA PEARL: It's done amazing thing because we didn't realize there are so many low-hanging fruits.

MARTIN FORD: Looking to the future, do you think that neural networks are going to be very important?

JUDEA PEARL: Neural networks and reinforcement learning will all be essential components when properly utilized in causal modeling.

MARTIN FORD: So, you think it might be a hybrid system that incorporates not just neural networks, but other ideas from other areas of AI?

JUDEA PEARL: Absolutely. Even today, people are building hybrid systems when you have sparse data. There's a limit, however, to how much you can extrapolate or interpolate sparse data if you want to get cause-effect relationships. Even if you have infinite data, you can't tell the difference between A causes B and B causes A.

MARTIN FORD: If someday we have strong AI, do you think that a machine could be conscious, and have some kind of inner experience like a human being?

JUDEA PEARL: Of course, every machine has an inner experience. A machine has to have a blueprint of some of its software; it could not have a total mapping of its software. That would violate Turing's halting problem.

It's feasible, however, to have a rough blueprint of some of its important connections and important modules. The machine would have to have some encoding of its abilities, of its beliefs, and of its goals and desires. That is doable. In some sense, a machine already has an inner self, and more so in the future. Having a blueprint of your environment, how you act on and react to the environment, and answering counterfactual questions amount to having an inner self. Thinking: What if I had done things differently? What if I wasn't in love? All this involves manipulating your inner self.

MARTIN FORD: Do you think machines could have emotional experiences, that a future system might feel happy, or might suffer in some way?

JUDEA PEARL: That reminds me of *The Emotion Machine*, a book by Marvin Minsky. He talks about how easy it is to program emotion. You have chemicals floating in

JUDEA PEARL

your body, and they have a purpose, of course. The chemical machine interferes with, and occasionally overrides the reasoning machine when urgencies develop. So, emotions are just a chemical priority-setting machine.

MARTIN FORD: I want to finish by asking you about some of the things that we should worry about as artificial intelligence progresses. Are there things we should be concerned about?

JUDEA PEARL: We have to worry about artificial intelligence. We have to understand what we build, and we have to understand that we are breeding a new species of intelligent animals.

At first, they are going to be domesticated, like our chickens and our dogs, but eventually, they will assume their own agency, and we have to be very cautious about this. I don't know how to be cautious without suppressing science and scientific curiosity. It's a difficult question, so I wouldn't want to enter into a debate about how we regulate AI research. But we should absolutely be cautious about the possibility that we are creating a new species of super-animals, or in the best case, a species of useful, but exploitable, human beings that do not demand legal rights or minimum wage.

ARCHITECTS OF INTELLIGENCE

JUDEA PEARL was born in Tel Aviv and is a graduate of the Technion-Israel Institute of Technology. He came to the United States for postgraduate work in 1960, and the following year he received a master's degree in electrical engineering from Newark College of Engineering, now New Jersey Institute of Technology. In 1965, he simultaneously received a master's degree in physics from Rutgers University and a PhD from the Brooklyn Polytechnic Institute, now Polytechnic Institute of New York University. Until 1969, he held research positions at RCA David Sarnoff Research Laboratories in Princeton, New Jersey and Electronic Memories, Inc. Hawthorne, California.

Judea joined the faculty of UCLA in 1969, where he is currently a professor of computer science and statistics and director of the Cognitive Systems Laboratory. He is known internationally for his contributions to artificial intelligence, human reasoning, and philosophy of science. He is the author of more than 450 scientific papers and three landmark books: Heuristics (1984), Probabilistic Reasoning (1988), and Causality (2000; 2009).

A member of the National Academy of Sciences, the National Academy of Engineering and a founding Fellow of the American Association for Artificial Intelligence, Judea is the recipient of numerous scientific prizes, including three awarded in 2011: the Association for Computing Machinery A.M. Turing Award for his fundamental contributions to artificial intelligence through the development of a calculus for probabilistic and causal reasoning, the David E. Rumelhart Prize for Contributions to the Theoretical Foundations of Human Cognition, and the Harvey Prize in Science and Technology from Technion—Israel Institute of Technology. Other honors include the 2001 London School of Economics Lakatos Award in Philosophy of Science for the best book in the philosophy of science, the 2003 ACM Allen Newell Award for “seminal contributions that extend to philosophy, psychology, medicine, statistics, econometrics, epidemiology and social science,” and the 2008 Benjamin Franklin Medal for Computer and Cognitive Science from the Franklin Institute.

JUDEA PEARL