Chat with Judea Pearl ACM Turing Award recipient in 2012 (Nobel Prize of Computing)

[intro music]

Stephen Ibaraki: Welcome today to our interview series with outstanding professionals. I'm Stephen Ibaraki, and I'm conducting an exclusive interview with Judea Pearl, ACM Turing Award recipient in 2012. And this award is really the Nobel Prize of computing. Judea is a legendary pioneer and world-renowned, distinguished researcher and professor at UCLA.

Judea, you just have an amazing background, with a lifetime of outstanding research contributions, just influencing so many diverse domains and with such lasting global significant impact. And we're really privileged to have you come in today, so thank you for sharing your considerable expertise, insights, and wisdom with our audience.

Dr. Judea Pearl: Thank you, Steve, for having me, and I'll be glad to share whatever I know with your audience.

Stephen: Judea, when did you hear of this extraordinary honor, and really, as I mentioned, the recipient of what is widely considered the Nobel Prize in computing, and this is the 2011 ACM Turing Award? And then how did you feel at the time, and what was the reaction from your colleagues and your family?

Dr. Pearl: We heard about it on the phone in the beginning of March. And to tell you the truth, I thought it a prank call from some Irish lottery, and I was about to tell them to put me on a do-not-call list. But my wife was on the phone. She said, "What kind of award?" [laughs]

And when she mentioned Turing, then I realized, yes, I may be on the list there, and I was very happy that the committee considered me. I was also partly surprised because I do not see myself a member of the mainstream computer industry. But, at the same time, my colleagues, when they heard about it, they all came out with the reaction, "What took them so long?"

So, a combination of pleasure and surprise, because my main focus of research is not considered mainstream computer science. It's not aimed to make a computer more powerful, more accessible. But it has a focus on artificial intelligence, which is a sort of, I would say, foster daughter of computer science.

And my focus is also on the cognitive science and on empirical sciences, which are at the periphery of computer science. I was happy to see that this jury, award committee, appreciates the impact of my work on both the AI and on peripheral themes around computer science.

Stephen: So what was the reaction of your colleagues and from your family?

Dr. Pearl: My family had a mixed reaction. My daughter asked, "Are you finally famous?" or "How is our life going to change?" and questions of that sort. But I assured them that nothing's going to change, I am going to remain the same daddy, and the award is just going to give me additional recognition and perhaps exposure so that my work will be appreciated, and I'll have a greater facility, a wider stage to propagate the ideas that I've been working for so long on. I think I pacified them.

Stephen: You definitely will get a wider stage on which to be recognized because of the stature of this award. Now, to the next question. What specific challenges in your education in Israel were catalysts or inflection points in your lifetime of contributions, and how and why did this happen?

Dr. Pearl: I never thought that my education was unique, that my upbringing was in any way special. But now that people keep on asking me about it, in retrospect I do think it was unique and special.

First, I was lucky to have giants for teachers. My teachers were all professors, academicians that came to Israel from Germany. They taught in academia in Heidelberg and Berlin, and they had to leave their positions and to take on high-school teaching positions. My generation was lucky enough to benefit from their expertise and their knowledge and grew up in a certain atmosphere of scientific challenge.

And they were polyglots, in the sense that each one of them could teach geography, history, math, and science at the same time, without notes, and we were just exposed to it by chance. And that had a great impact, I think, on my education. I'm speaking in retrospect. At the time, I didn't realize that other educational systems do not have this luxury, that privilege.

And the second aspect of it which was unique was that my education was embedded in an overarching cultural renaissance of nation-building, with total commitment of the entire society to excellence and a collective ideal. Everything was in the making. All molds are to be broken, and the future is what we make out of it.

So, evidentially, this spirit percolated into science education. I remember one of my math teachers gave us the feeling that you, too, can make an original contribution to math. You, too, can find an original proof of Pythagoras' theorem, for example, or [laughs] a proof that has never been thought of before, or a new slant on Newtonian mechanics, one that Newton perhaps had neglected to consider.

He even had a seminar in the summertime, on his own time, with a few kids, enticing them for original work in mathematics. And the same thing applies to college. We had great teachers that always spoke as if they were present at the time of the discovery. They made the discovery vivid, in historical context, and vivid to our minds. And we were spoiled enough to assume that we need to understand the lectures in real time. That's something which I have not seen at UCLA in my teaching.

In American education, students are more or less passive. They do not interact in real time with the professor. They tell themselves, "It's OK if I don't understand the lecture at the moment. I'll take the time enough to learn to go over the notes, and perhaps I'll understand later." This is unheard of in the Israel.

You've got to understand in real time, and if you don't understand, you let the teacher know, either by raising your hand or moving your chair or resorting to other kind of mischievous activities. Understanding in real time was considered to be standard.

So I was lucky in that sense. And also, I should say that I'm lucky that higher education in Israel did not suffer, though I took my college education during the time of austerity. We shouldn't forget that in the 1950s, Israel tripled its population, threefold, from 600,000 to 1.8 million, and no one went hungry, and higher education did not suffer any budget cuts.

And that was due to the great vision of the leadership, who understand that if you build a nation, scientific excellence is a prerequisite. And we were simply the beneficiary of that, I think, vision and leadership. I'm constantly reminding myself and my audience that I owe a lot of what I've done to my education and to the spirit of nation-building in Israel.

Stephen: That's a really interesting perspective. And you talk about the fact that, at that time, even though it was the early days and it was a time of austerity, there was no budget cuts, and there was passion in the classroom and excitement and interaction. How would you contrast that with the current education system in the United States?

Dr. Pearl: I think it has to be revamped. One of my pitches is that mathematics and science, in general, ought to be taught in the way it was discovered, in historical context, not hierarchical in terms of the logic of the subject matter but in the way it was discovered, with the scientist's picture on the wall, even the scientist's family on the wall, I'm not kidding, because this is part of the context of the discovery.

Archimedes in the bathtub. It's important. And that gives students the idea that science is not a dry subject, but it's a human subject. It's a description. It's part of the struggle of the human mind with the mysteries of nature. And that means that the student becomes an agent.

Stephen: And I guess you would have that same sort of feeling about budgets. You would definitely want to ensure the education system in your current country, and that is the US, that budgets are increased, I guess, right?

Dr. Pearl: But it's not only budget. It's good teaching, teaching from the heart, combined with budget, because budget only enables you to train and to attract good teachers.

Stephen: Now, I asked you about some of the aspects of your education in Israel and how they were catalysts or inflection points in your lifetime. Let's apply that same question to your work at RCA and then your move into academia. How did your work there become catalysts or inflection points in your lifetime of contributions, and how and why did this happen?

Dr. Pearl: My first job was in RCA research laboratories in Princeton, New Jersey. And during the '60, that was my first exposure to computer science, or to computers. And everybody was ecstatic by the potential growth of the field. Everything was new. Every phenomena in physics you can think of, referred to a potential to becoming the future technology for computers.

So everyone understood that the days of vacuum tubes and even the transistors are numbered and were looking for new phenomena, new physical devices, new physical phenomena to implement computations.

Memories were a stumbling block. At that time we used magnetic-core memories, which were donuts that literally thousands of girls around the globe, from Hong Kong to Massachusetts, had to string, one by one, crisscrossing the envelope. And everyone knew that the days of core memories are numbered, and people worked on all kind of phenomena.

Some worked on new magnetic configurations, and some worked on photo-chromic memory. And I worked on superconductivity, because in superconductivity you can store information in the form of turbulent currents going around, like vortices, without needing to string them. And that was an exciting time.

I remember our supervisor, at least once a week he would get into the office and say, "Well, we have a new patent disclosure." The whole atmosphere was rushing and chasing after new patents. [laughs] However, eventually, all these activities came to a halt because semiconductors took over. Everyone who worked on memory devices had to find new jobs.

Nothing remained of my PhD dissertation, which had to do with superconducting vortices, except for one piece of analysis which I did, which recently I noticed physicists without my knowing called the "Pearl vortex." It gave me an instant claim on immortality.

But at the time I had to find a new job, and I went to academia. It was easy then to go from industry to academia because industry looked down at academia. Nothing of importance came out of academia. The transistor and the laser, they were invented in industry.

So I was accepted to USC as a memory expert, but with no activity in memories or in the entire computer field on memory devices. Everything went toward the semiconductor memories. And that activity was taking place primarily in industry.

[inaudible 00:15:48] camera was pioneered at the time, so I went to work on the things which interested me privately for a long time, namely emulating myself, and that's something which I think every computer scientist is curious about, more so than psychologists, because computer scientists do have the means to emulate themselves.

And I started working on pattern recognition and various aspects of vision. But eventually I ended up with heuristic search and probabilistic reason, causality, and so forth. And I got into artificial intelligence research.

Stephen: You know, you've made so many contributions, so how has your work specifically and lastingly contributed to...and there's various components here. things like machine reasoning and natural language processing and computer vision, robotics, computational biology, econometrics, cognitive science, statistics, philosophy, psychology, epidemiology in the social sciences?

Dr. Pearl: Well, the list may give the impression that I'm expert in all these fields, which would be wrong. But if you notice, all these fields that you mention have one thing in common, and this is uncertainty. They are embedded in uncertain and noisy data, and they all need to have principles of filtering out the noise and extracting out meaning. This is where my work on uncertainty comes in.

I should then divide it into several categories. I believed for a long time that probability is a dominant formalism for handling uncertainty. And the time we were contenders to probability theory. I'm talking about the time, the earlier days of artificial intelligence, especially the early days where expert systems came into being. What people used to do is take the deterministic rule based systems and decorate them with uncertainty or degree of belief.

If you have a rule, for instance, that if you have a slight fever, that you are likely to have malaria. A rule like that will be quantified with a certain degree of strength. And then these degrees of strength would interact when you have a large expert system with many rules, until you get the right answer, which is solidified and condensed the uncertainties giving to all the rules in some mysterious way, which was not clear even to the programmers.

And the question, at that time, I strongly believed that when you use probability theory, you are guaranteed to have intuitively acceptable answers to your questions, because in textbook probability, you rarely get a surprise. You rarely get a paradoxical result.

There are some paradoxical cases or in probability theory, but very few. Most of the time if you put in the right data, if you put in the right rules and knowledge, then you expect to have an intuitively acceptable or plausible answer.

So I asked myself if probability is the protector of plausibility, let's put it onto a computer. The trouble was only that even if you have a few variables, the amount of storage that you need to invest to encode a probability function for 10 variables, it grows exponentially. The question was if we do it probabilistically, how can computers cannot do it probabilistically?

And if we do it, despite the fact that it calls for exponential storage, it means that some approximation is feasible. From the fact that we find probabilistic answer to be plausible, and we do things fast and in a lively, we have a proof that approximation is feasible. And if approximation is feasible, let's put the approximation on a computer.

And I looked for a long time for such an approximation. It became clear that if you break down the problem into loosely coupled sub-problems, you can get parsimonious encoding of probability knowledge, and you can get tractable algorithms for processing it and for updating it.

And you also get a way of debugging it, mainly if things don't go according to expectations, you can find out which rule is the culprit and debug it. This combination of tractability, storage economy, and the re-configurability that was exactly what is needed to make, to change expert systems from rule based to probabilistic.

And, at the same time, it gives us something which is very called for by cognitive observation, cognitive science observation. And explain how people are able to read the text fast, to park their [inaudible 00:22:18] reliably and perform operations in an environment which is constantly bombarded with noise and erratic phenomena and partial knowledge, and do it so reliably and swiftly.

So it gave us an idea of how people do things, and it also gave us the hope of being able to do it distributively. And distributively is important, because it's hard to believe it in our mind we have a graph supervisor. Let me...an algorithm to tell each neuron when to fire. Evidently, neurons are firing in response to their neighboring neurons. So they're firing in their neighboring neurons.

And their things are done asynchronously, autonomously. Each neuron can wake up in the middle of the night and decide I want to fire now. And the system should be able to accept, to propagate all these firings and reach a certain equilibrium, at the end of which, each concept is given the correct degree of belief that is dictated by, by what? By probability theory! Otherwise, it wouldn't be plausible.

That gave us the impetus, the drive, to look for a structure that could enable to propagate belief asynchronously. [inaudible 00:23:58] went into trees and poly-trees and then other structures, but eventually there was a dilemma, what do you do when you have loops and you have a constant...you have a continuous propagation that does not convert, does not convert any of it's converse, converts [inaudible 00:24:24] to the wrong answer.

Things of that nature kept on creeping in this work, and eventually out of desperation, I suggested that the mind probably neglects the presence of loops and instructs each neuron or each module or each variable to be activated and to send messages as if it was part of a singly connected network. Namely, as if it was part of a loopless network, a network miraculously will. It changed the whole ball game and gave...and produced a result that are much faster than what one would theoretically anticipate.

And it became a tool, a software tool, for anybody who is dealing with uncertainty to be able to update beliefs in response to observations and in response to different knowledge, to updated knowledge, and to get the correct belief, reveal at reasonable times and using a reasonable amount of storage.

That explains why the beneficiaries of these techniques are so widely spread, from machine learning to natural language processing, vision, robotics, and so on. All these areas must deal with noisy environment, noisy observations, noisy data.

The second step which one should single out is the work on causality, which was my second transition. I left probability theory, and I went to causality because I realized that people do not store knowledge in the form of probabilistic relationship, but in the form of cause effect relationship.

And that all this probability propagation and so on. And all the beliefs and all the arguments which we invented for belief propagation are just the source of phenomena. And they account from cause effect relationships.

So I went to...I shifted, totally, from probability to causality. And the two, by the way, do not mix. We have a different language, like the saying goes, "That correlation does not imply causation." And that is revealed in the language that one needs to use to handle these two fields of relationships.

The hardest thing for me at that time was to quit probability as a language and talk about computer language, OK? To quit probability as the Oracle of knowledge and to move to a new language, which is a Cause/Effect language when we deal with causality, when we deal with intervention, when we deal with action.

That was a real, I think, dramatic shift for me, and not only for me, for...it still remains a dramatic transformation for most scientists, because we all think statistics 101 and probability 101, and we are led to believe that if you had probabilistic knowledge, then all the answers in the world...all the questions in the world would be answered. And this is simply wrong.

Causality demands a language of its own. And once you realize that, then as a computer scientist, you can invent the language if you don't find it in the literature, you invent it. We take linguistics shifts from one language to another with a smile, with grace, and most scientists take such a shift to be traumatic or impossible. I mean, I keep on saying that people go to war for languages more than for the content of the language.

And, anyhow, we are talking about causality now so the area of robotics, the areas of cognitive science, definitely needs an understanding of how people process cause/effect relationships. If you want your robot to be able to handle uncertain environments and things don't go according to anticipation. Like if the robot spills the milk or breaks the glass then that robot need to be able to repair the environment, and such repair can only be done if the robot understands the cause/effect of the relationship in that environment.

So the question comes how does the robot learn cause/effect relationships, and how the robot then utilizes that relationship once it is learned or once it is even taught by the programmer. It's not an easy task because the rules of cause/effect are not logical rules. These are new...they obey a new calculus, a different kind of calculus with ordinary and logical implications.

So that is exciting. We do things in our minds and we don't have a calculus for doing it. This is a dilemma which is enticing to every computer scientist. I think. At least it was for me. So I spent [laughs] a lot of time trying to formulate such a calculus that will enable us to handle actions and observations and come up with a correct answer to a question, for example, what will happen if I intervene?

Once we understand why it is important for cognitive science and why it is important for robots, we can also understand why it is important for all the empirical scientists without working day and night around computer science. And I'm talking here about economists. I'm talking about statisticians. I'm talking about psychologists, [inaudible 00:31:28] , and social scientists.

These are hardworking people spending millions and billions of dollars on randomized experiments, trying to figure out if a drug has beneficial effects on society or adverse effects, and trying to find out what will happen if he will use class sizes, what will be the quality of education?

All this requires careful experiment and a lot of money is spent on that, and if we're all immersed in a noisy environment, tons of data, and they can be likened to the robot in the middle of a kitchen or a laboratory trying to make sense of the signals he or she receives.

So I liken them all these hundreds of thousands of empirical scientists to a colony of robots. From a cognitive viewpoint they are facing exactly the problem that a robot is facing in a new and uncertain environment. And they do not have the principles for handling that environment. They're simply working with traditional probability and statistics, and these languages are not satisfying their needs.

So I was lucky, and we were lucky, that the effort we spent on trying to equip robots with the capability of reasoning with cause and effect enabled us to unveil the principle by which human scientists should analyze cause/effect relationships and we were able to equip them with a powerful calculus for doing that.

And it was important because throughout my research in robot understanding, I had the support of hundreds of thousands of customers in the peripheral field of empirical scientists who were thirsty for any new trick, any new algorithm, any [laughs] new innovation in the study of cause/effect relationship.

And that explains, I think, why the work I'm doing happened to be applicable to so many fields. We are dealing here with the basic of reasoning under uncertainty and under cause/effect relationships.

Stephen: And then I guess the causation kind of research that you did really ties into philosophy itself, right?

Dr. Pearl: It tied to philosophy because philosophers for a long time tried to explicate what cause/effect is, and it...actually it's a nice story that most of philosophers who dealt with this question referred to counterfactuals as the basis for explicating cause/effect. [inaudible 00:34:45] , for instance, in the 18th century, defined cause as in terms of counterfactuals. A is said to cause B if B would not occur had A not occurred.

We went background and reasoned against backward in time counterfactually. We know that A and B did occur. When do I say that A is the cause of B? If B would not have occurred had A not occurred. So you see here we have a giant of philosophy, a pillar of philosophy, explicating cause in terms of counterfactuals. And the same thing with David Lewis, which is a modern philosopher also basing cause/effect on counterfactuals.

And the question is why did they explain causation in terms of counterfactuals and not the other way around? It must be that counterfactuals are more cognitively plausible, more...less problematic, more deeply entrenched in our intuition, which means that people can process counterfactuals more swiftly and reliably and consistently than they can process cause/effect relationships.

Wow, if this is the case, then why shouldn't we equip machines with counterfactual reasoning, and why shouldn't we base causal reasoning in machines on the basis of counterfactual knowledge? And indeed I believe today, and I realize for maybe [inaudible 00:36:30] that counterfactuals are the building blocks of scientific thinking.

There couldn't be any science without counterfactual. The whole physics is one big [inaudible 00:36:42] of counterfactual reasoning. The equations of physics are not algebraic equations.

They're counterfactual equations. If I give you a simple rule, if the weight is, let's say, Hooke's Law, OK, "the length of the string is proportional to the weight that you hang on it," it operates under any weight.

In addition to the current weight, which happened to be one kilogram, the law tells you what will happen if the weight was two kilograms, and three kilograms, which means every physical equation that we find in high school physics really stands for a collection of counterfactual statements. If the current would be three amperes, the voltage will be that much. Every law in physics, right, can be interpreted counterfactually.

And the conclusion is that we need to build a calculus for counterfactuals, and also the conclusion is that if you encode counterfactuals in the form of physical equations, here I'm talking about structural equations, although the principle is there is a dilemma. How can we human beings communicate with counterfactuals and understand each other so reliably and so consistently, and we form a consensus?

For instance, if I tell you that if Oswald did not kill Kennedy, then someone else did, you would say, "Yes, that makes sense." And if I tell you that had Oswald not killed Kennedy then someone else would have, you say, "Ah, I'm not too sure. I don't think so." And this is uniformly understood by people in our society, and we form a consensus about counterfactual things.

Even if reason backwards to Julius Caesar we can think counterfactually what will happen had he not crossed the Rubicon? That's how we teach history. And we understand each other. And if we do it means we have a counterfactual machinery in our minds, and we store knowledge in such a form that it enables the machinery to operate and deliver answers and draw conclusions swiftly and even reliably.

Now if this is the case then we should do the same to machines. We should be able to. And the fact that we have an existence proof in the human skull gives you the encouragement to commit yourself to a piece of research and believe that it's going to be successful. It must be successful. In the case of counterfactuals I'm very satisfied with the result that we have.

We do understand counterfactuals today, and we have been successful in algorithmitizing, of reducing counterfactuals with into algorithms, and that is a tremendous ramification both for robots and for empirical science, because then we can questions for instance, about credit and blame.

For instance, is this employer guilty of sex discrimination in hiring? This is counterfactual which had he not discriminated, this lady would have been hired. There are many cases in the law which are defined counterfactually, and a lot of questions in the empirical sciences have a counterfactual flavor.

So I'm very happy with the level of understanding that has been reached now. And again, reached by what? By an attempt to equip robots with that capability. The beneficiaries are all the empirical sciences that rely on counterfactual thinking, and they all do.

Stephen: Can you additionally profile your extensive research history, their lasting impact, and some valuable lessons you wish to share from each of your top research areas?

Dr. Pearl: Mm, valuable lessons. I can summarize my valuable lessons in three sentences. First, people reason with probabilities. And second, people do not reason with probabilities. But people [laughs] reason with cause/effect relationships. The third is that people do not reason with cause/effect relationships. They reason with counterfactuals.

So here you have the three sentences, the major impact, and lasting impact of my research. The impact itself is in the algorithms that were developed to encode these relationships, to produce new influences from observations and from knowledge. But one should classify the impact of my work into three layers: probability, causality, and counterfactuals.

Stephen: I can see how your progression, and your work, and your research, and as you gained insights, and sort of your insights evolved, that become sort of the major inflection points in your life, but also the major reflection points in your field as well.

Dr. Pearl: It has had impact on the field. It has given researchers in AI the comfort and the confidence that they can encode knowledge in this form and can perform inferences on that knowledge, and also it has a tremendous impact, as I explained, on peripheral fields.

One area, for instance, which benefited immensely, one question that has been lingering on for generations without any even attempt to a solution is the question of external validity. Let me explain. You conduct experiments in the laboratory, and you assume that the results will be applicable to outside the laboratory, to the field, to the population.

A robot is being trained in a simulating environment, you expect the robot to apply what it learned in the simulation room outside in the field. So, under what condition can we extrapolate from the environment in which we were trained, in which we learned, into a new environment which differs from the first in several aspects?

Even if you know the aspects which make the two environments similar and different, how would you transfer the knowledge? This has been an unsolved problem first posed by empirical scientists. It's called external validity. Others call it extrapolability, generalizability and other names, but it hasn't been solved. Not even touched.

We found a complete solution for that problem, which means that if you are able to encode the aspects of the two environments that you believe are common, and those that you believe are different, and you tell us the list of differences and commonalities, we are able to tell you whether a certain relationship that was learned in the experimental environment is transferable, with all the external knowledge that you can have, into the new environment, where you cannot conduct any experiments, but you have to apply the knowledge that you learned in your training.

I am very excited about it. It's a new result, and I'm waiting for all the social scientists, and economists and epidemiologists to jump in glee and to rejoice into this discovery. Well, it's a year already. They haven't jumped and they haven't danced, but I have patience. I know that one day they will.

So, this is an example of a problem that was handled well by virtue of what? By virtue of our ability to take an unformulated question, an unformulated system of relationship, and submit it into a formidable and draw the conclusion, and algorithmitize it. The same goes to, the application for robotics is also clear.

If we can handle counterfactuals, then we can handle retrospective and introspective thinking, and we can equip robots with what we call free will, and responsibility. That will be a great advantage for robotics and for human-machine interaction, because you and I communicate under the assumption that there is free will, and we transfer knowledge with that assumption. We make each other do things with that assumption.

Equipping robots with such an illusion will improve the communication between men and machine. It will also teach us a lot about our cognitive system. What makes us obtain? Under what conditions do we have the illusion that we act out of free will, and when do we have a different sensation? If we act involuntarily.

We may also be able to answer the question of what we gain by this communication language. How would a soccer team of robots play better equipped with that facility as opposed to the way robots communicate today without the facility to assume social responsibility?

Those are partly philosophical, partly psychological, but mostly engineering problems. That is the nice thing about it. You take a philosophical loose problem, fuzzy problem, decorative problem and we reduce it to an engineering problem.

When we assume engineering content, you cannot fool things around. You cannot slow people. You've got to show performance. When you harness logic and the principle of computer science semantics and synthetics and you demonstrate performance.

Stephen: You mentioned the sole concept of robotics, and it's clear to me that from your discussion there that you do believe that we will have robots where free will can be developed because it's really an illusion, and you can engineer that illusion using counterfactuals. Is some of that work being done, then, at some of the laboratories around the world that use your research and encode it within a robot?

Dr. Pearl: In a certain way, you can say it's done everywhere. You can say for instance that the robot that plays chess is living counterfactually because it investigates a variety of moves. So it looks at several moves, and chooses one that looks best according to a certain criterion. But, because it did explore other avenues, you can say it, in some loose and superficial way, reasoned counterfactually.

When you ask the question, "Would robots be able to be equipped with free will?" one could argue that any chess playing machine already has free will. Otherwise it wouldn't have chosen the move that it chose. So it's a superficial way of thinking, but it's a very deep, philosophical question now. Under what condition can we say that the robot has free will?

It's like passing a mini Turing test. A Turing test for free will. When we talk about human beings, the arguments goes immediately philosophical. Whether we have illusion of free will, or truly have free will.

Most scientists that I talk to believe that the free will is an illusion. Some believe that it's not an illusion, but belies all kinds of mysterious phenomenon such as the Heisenberg Uncertainty Principle, who knows. Fractals, everything that they can hang on to that can explain how we have the strange sensation that we have the option to do one thing, or not to do it.

At the same time, we have this very clear sensation that we can do, or not do certain actions, when in fact we do not have this option. We are going to act one way or the other depending on how our neurons are firing at any given time. So, there is a conflict here.

When we talk about human beings, the conversation immediately becomes philosophical and unmanageable. But when you talk about robots, it becomes clear because no one would claim that robots have free will.

So then the question is, can we equip robots with free will? It becomes an engineering problem. Under what condition would you not be able to? You would not be able to distinguish between a robot or an organism that has free will, and one that does not like current day robots.

I think the solution lies in making a robot that can explicitly speak about his or her will. Like, why did you want to do that? Once the robot is equipped with that ability, then you wouldn't be able to distinguish, and according to the Turing test, a robot that acts as though it has a free will, does have a free will. That's not a question anymore.

With that comes, for instance, social responsibility. Social adaptation. An agency that all comes with the idea that, and communication through credit and blame. One robot should be able to tell another, you acted wrongly, and you should be blamed. You should be put out of our team. You caused us to lose the game.

The other robot, the blamed one, should be able to take that communication and translate it into a reconfiguration of its software. That would be then the proof that the robot has free will, and he or she understands the meaning of credit and blame.

Stephen: It's interesting. Earlier you talked about this whole idea of external validation or validity. Is that really where your current research is focused? In other words, can you define what would be the outcomes of your current research and their applications?

Dr. Pearl: I cannot define, I can only hope that it will lead to a more friendly conversation between man and machine. It will lead to a more effective communication among robots when it comes to issues such as social good, collective aspirations.

All these concepts that move us into action. If we can transfer them to robot communities, that will be a great achievement in terms of the ease with which we can communicate with robots. Because we are social animals, and in order to communicate with robots, we have to be able to equip robots with understanding of social interactions. This counterfactual is a basic building block in social interaction.

So this is one aspect where I see the results of this research developing into a useful software. External validity, that will happen not in the robot world, but in the empirical sciences. Of course, it will affect machine learning, because within machine learning, you don't want to relearn things as you move from one environment to another.

You want to use all the aspects of the training environment that are shared by the new environment, and just learn the novelty in the new environment. Otherwise relearning everything from scratch is a lost cause. It's a losing game. That's where external validity comes into game. It will improve machine learning tremendously. I don't know how at the moment, but I just know that it will.

Stephen: But it's clear that from the applications of external validity, and really it hasn't been taken up as broadly as it should be, or it's going, but that could lead to a Nobel Prize in economics.

Dr. Pearl: The economic [inaudible 00:56:16] that I know are very conservative. They are the last ones to utilize something that they can benefit from. So I think it will be picked up by machine learning first, and then transferred by osmosis to more conservative circles, among them social science, economics, health research.

That brings us to the question of how the Turing Award may help me propagate the ideas into laboring fields. So I'm extremely hopefully that a new stage is ahead now, and a new microphone, that will enable me to propagate this idea to people who can immediately benefit from that.

Robotics people will benefit from them in 10-15 years, once we have a community of robots playing soccer. Not today. But the people who are dealing with economics research, the people who are dealing with epidemiology, people who are dealing with education research and psychology, these can benefit tomorrow. So this is one hope that I have for the spreading of these ideas.

Stephen: Interesting. So what are your most difficult challenges in research, and what valuable lessons do you wish to share from these challenges?

Dr. Pearl: Well, my greatest challenge was the transformation from one language to another. From probability to causal, and from causal to [inaudible 00:58:03] . And I know that if I had to teach young people from my experience, the thing that I would encourage them most would be to have a flexibility for shifting languages.

I found this to be the hardest obstacle in my research. They were the hardest obstacle to overcome in my research. I see them being the hardest obstacle to overcome among my colleagues, especially the non-computer science colleagues. And I wish we could revamp education to the extent where growing scientists will have the flexibility to shift from one language to another.

Stephen: How do you create that ability to shift from one language to another?

Dr. Pearl: I think by demonstrating. By demonstrating that you cannot solve certain class of problems within one language, and you can solve them easy with another one. And demonstrate the ease of transformation.

If you equipped a language with one or two new operations, with one or two [inaudible 00:59:38] for handling the new operations, then “zoop!” All the questions that you couldn't answer in the first language are answerable swiftly and reliably.

That has been the greatest pleasure that I drew from my research, and I believe it can be transferred to students in every field. In every field that is cattled with the idea that you can tackle an area which is [inaudible 01:00:15] but they had no prior knowledge about.

And once you have the vital language, a light shines and you are able to handle difficult questions with the new tools. That is the greatest pleasure that I drew from my research. Being able to do intellectually today something which I couldn't do yesterday, simply by learning another trick or another operator. That is a tremendous pleasure.

It's like growing up again. You can do things today that you couldn't do yesterday. By doing, I mean not just multiplying larger numbers, I mean tackling questions that you didn't know how to even approach. You can tackle them today with ease and comfort, and smile. That is the greatest pleasure I drew from my research.

And I saw it happening. Questions I couldn't even begin to approach became easily solvable once I get through certain transformation.

Stephen: I can see that your work in external validity, and people accepting that work and shifting to this...all the possibility that this enables. But there's always the idea of group think, where people are resistant to change. Why do you think that's so?

Dr. Pearl: Well, there's many components. I focused on the transformation from one language to another. Indeed, language is the greatest guardian of dogmatism, or propagator of dogmatism.

People learn languages early in their childhood, and then they cannot shift. And that happens in science, too. When you learn one language, it's hard for you to shift, and you're getting to this cult.

It's really a cult of believing that if you only knew, if you only were able to express the knowledge that you have in the language that you have, then you'll be able to solve things given that the language you have is incapable of expressing what you want it to express.

That I can see again and again among my friends in for instance, in econometrics. Economics, for instance. If you look today into the definition of econometric models on Wikipedia, you can find the definition which is simply wrong, and reflects exactly what we were talking about. The inability of econometricians to understand that economical relationship cannot be expressed in the language of [inaudible 01:03:12] .

So definition they give for econometric model is...it is a collection of rules...it is a statistical model which captures economic relationship. Which is dead wrong. It is not statistical model. It goes beyond statistics, you cannot express with statistics the content of economic model.

Like the effect of let's say duties on prices. You cannot express it in the language of statistics, because it has to do with the word effect. Effect is cause effect relationship, and statistics is incapable of expressing things like as simple things as the rooster crow does not cause the sun to rise.

It's impossible to express that in the language of statistics. And econometricians have not learned that, have not extract...be able to liberate themselves from this dogmatic language, to believe that statistics can answer all their questions. This is an example of conservatism that is induced by sticking to a language, that inability to change language, I will use an example.

Plus there is also inertia in education and in the research community, and in I say politics. The politics surrounding the transformation of paradigms. People have power over other people, and influence other people to use one vocabulary as opposed to another in order to propagate their ideas or to resist other ideas from spoiling the youth. But I think the main one is the ability to liberate yourself from the constraints of language.

Stephen: We're discussing some of the challenges in terms of transitioning from one kind of language to another kind of language and opening up a whole bunch of possibilities. Can you describe other areas which you may consider controversial or where there's much discussion in areas that you research?

Dr. Pearl: I can discuss, for instance, there is a great controversy now in statistics, whether counterfactual reasoning is beneficial. You cannot measure a counterfact. You cannot refute a counterfactual statement because it is introspective about things that didn't happen. What would have happened had I not done what I actually did?

So these are philosophical questions. And the great controversy is whether all statisticians should resort to such a language. Again, I'm talking about the language. But for them it's not a matter of language, it's a matter of scientific true and false. If you cannot test something, that it is meaningless. And if it is meaningless, then it might lead to erroneous results without your knowing.

That is one controversy which I am immersed in the midst. And I do not believe in that, by the way. For instance, we have many things in mathematics which are a fiction of our imagination, complex numbers, to mention an example. Or even negative numbers are not observable. They're only placeholders for some arithmetic operation in our minds or in our algorithms.

So coming from this viewpoint, I am for using, in ordinary research and ordinary inferences, concepts which are not directly testable, which are testable only indirectly, in terms of the results which they produce.

And if the results can be proven to be always in conformity with things that we do measure, then we shouldn't hesitate to resort to such fictitious entities that mathematics endows us with. And counterfactuals are such entities. Very useful and highly controversial.

Stephen: I guess another area would be this whole idea of religion and religious myths being metaphors or poetry for genuine ideas that are difficult to express. Do you want to talk about that at all?

Dr. Pearl: Yeah. Many scientists get irritated by religious talk. Given that they are atheists and they do not believe in a deity or in life after death or in supervision over one's actions, they get irritated when people do communicate in such a language, like inter-religious, interfaith dialogs, which are truly meaningless if you are a dogmatic atheist.

I'm an atheist. I don't believe in life after death, and I don't believe in higher supervision over one's action. But at the same time, as computer scientist, it behooves me to translate human communication into an algorithmically meaningful context. And I find the communication in nonbelievers to be meaningful. It's loaded with metaphors, such as you'll be punished, you will be burned in hell, [laughs] you're condemned, you'll be redeemed, and so on.

On one level, it's meaningless. If you don't believe in life after death, what does it mean you will be redeemed, or you accept Jesus, whatever? But at an algorithmic level, it makes a lot of sense. It's a metaphor for some punishment that is realizable. If not after death, it's realizable as a collective in the course of history.

The whole idea of ethics means that you are subjecting your instant gratification into a higher good, and higher good means into a collective good in the long run. So from a historical perspective, all this religious metaphor makes sense. "God will punish you" makes perfect sense.

If you replace the word "God" with the word "history," suddenly the whole Bible makes sense. Which means, in the long run, history will punish you. Which means, if your evil deeds will be picked up by young people and they will perpetuate your evil behavior, eventually society will suffer. History will suffer. You, as a member of the collective, will suffer in the long run. "God will punish you," now it makes perfect sense.

So if you interpret religious metaphors in this context, I open the radio and I listen to an evangelist preacher, he makes perfect sense to me. Even though I'm atheist, I'm Jewish, but I understand him perfectly. And that gives me some, I think, power in communicating with people from different faiths.

I don't laugh at what they say. I respect what they say because I've translated them to a new computer language and I understand the meaning of what they're trying to convey. It's as simple as that.

So the question comes, of course: should we equip robots with religion? I don't know the answer. It will be a nice experiment, to see if a community of robots becomes more friendly or socially responsible once we equip them with the illusion that there is a deity that overlooks their deeds and their thoughts and takes records. I don't know.

They may become vicious like some of the religious cults or they become socially exemplary social animals. I don't know the answer. The only experiment we have is the experiment with human beings, and we are still debating whether it was successful or not.

Stephen: Judea, you've laid many of the foundational pillars as one of the top groundbreaking visionary innovators and distilling from your experiences, what are the greater burning challenges in research problems for today's youth to solve, to inspire them to get into computing?

Dr. Pearl: I think the greatest advice I could give to youth would be to listen to their intuition and to listen to their century-old curiosity to understand themselves. And this is I believe the driving force behind most computer scientists, an urge to emulate yourself, but first it's an urge to understand yourself.

But today we learn that understanding implies ability to emulate. When you emulate, you can take things apart and you understand, and understanding means being able to taking things apart or to predict behavior under the conditions when things are taken apart.

So I would just advise them to listen to their intuition, intuitive drive to understand themselves and not to take no for an answer, because every aspect of human behavior that is not yet but can be implemented on a machine is a PhD thesis, and deserve a PhD thesis. So I would just advice them to follow their natural intuition in terms of understanding themselves.

Stephen: Now to our youth again, with an interest in the future of computing business, but without the educational foundation, how would you explain your work in probability causation, the calculus of intervention, fascia networks, and contrafactuals?

Dr. Pearl: Oh, this is really easy. I just need about 10 minutes in front of students and I can explain it. So everybody understands that things are uncertain in the world. There are all kind of noises. You make estimations, and the things don't turn out what you expected. That's uncertainty.

We humans are able to handle that fairly well. We cross the road, and we don't get hit by a car even though our estimate of the distance of the car wasn't very accurate. And in everyday activity we are bombarded with random phenomena and we are able to manage. The question is, how can computers manage the same degree of uncertainty?

And we learn in school that computers work by ones and zeroes, truth and false. How do you equip a machine with the degree, the shades of belief, the shades of likelihood? This is the work of probability. It's one layer, easy to explain. When students like to emulate it on the...with the three variables they get a great kick from seeing that computers can produce results which are plausible.

Give them a simple detective story like, for instance, Watson and Sherlock Holmes, and Watson...I have it in my book that Holmes gets a phone call from the neighbor saying that he heard a burglar alarm, and Watson is worried that a burglary took place, and then he listened to the radio and finds that an earthquake took place, and he wondered whether it was really a burglar alarm was triggered by a burglary or from an earthquake and so and so on.

Give them a story like that, and you say, "OK, encode the story on a computer and have the computer answer a few reasonable questions such as, 'Would the radio announcement about earthquakes reduce or increase the probability that a burglary took place in Watson's apartment?' And the computer answered nicely. It reduced it by 20 percent, which makes sense, because if you explain one cause, you take away belief in the other cause. It makes sense.

So students, even not computer scientists, get some of these kicks from seeing their knowledge emulated and amplified by the computer. And then they appreciate the difficulty of encoding more than three, four variables. They very quickly run into exponential exclusion, and they appreciate the ability to accommodate thousands of variables.

Then comes the question of cause and effect. If you just ask them to encode a simple sentence, which I mentioned before that the rooster crow does not cause the sun to rise, they appreciate the difficulty of encoding cause and effect relationships. You cannot encode it with the languages that students in any of the sciences have been trained. It's undoable.

So when you pose to them the same problem, encode the story about Watson and Sherlock Holmes on the computer and answer a few questions, but now the questions are going to be more sophisticated such as, "What will happen if Watson decides to do something and to, for instance, play a prank call to Sherlock Holmes and tell him a certain joke?"

Maybe if you ask a question that involves intervention actions they cannot do it with the available machinery that they have. And again, a tremendous kick for me being able to do it once you show them how to encode the idea of action. Do X. Raise taxes. Make me laugh.

All these are actions that can be encoded nicely on a computer, and then you ask a simple question, "What will happen if I do, or what will happen if I do things after I observe something?" and you get the right answer. What do I mean by right? One that matches our intuitive expectations.

Students have a tremendous pleasure from playing this game, and they also get appreciation for the kind of...for the general strategy of taking an area that has not been encoded before and encoding it and seeing how your intuition gets amplified by formal language and formal algorithms, yeah.

This amplification is...I think it's universally appealing to all students, for all young people, regardless if they are in computer science or outside computer science.

Stephen: So upon reflection, what specific qualities make you excel, and why?

Dr. Pearl: I'm the last one to answer this question, aren't I? I think that the answer should be...I'm lazy. I'm lazy and lazy. And lazy means that I cannot think very deeply, and I would like to have a pocket calculator that can think for me, and I found [laughs] one.

So this desire to transfer my thoughts due to laziness to a machine that can do things for me perhaps is...you talk about quality. That's a quality that haunts me throughout my life, being lazy. At the same time I should say, "Hard work, hard work, hard work." So I can give you two answers. Lazy, lazy, lazy, and hard work, hard work, hard work.

Stephen: Mm, that's interesting. Now looking at your past, present, and maybe into the future, name three or more who inspire you, and why is this so?

Dr. Pearl: You mean people, right?

Stephen: It doesn't have to be people.

Dr. Pearl: To tell the truth, I was very influenced by stories about great scientists. I was. I did run out of the bathtub together with Archimedes screaming, "Eureka, eureka!" I did. And I climbed with Galileo on the Tower of Pisa, and insisted to see the two rocks falling. And I did dance with Faraday. I did, when the two magnets started rotating when he made the first rotating machine.

So I think stories about great scientists did influence me a lot. And present day scientists, if I were to choose, it would be Boole and Shannon. Boole because he took...he translated logic, which before him was not a symbolic province. It was done in the mind, and he reduced it to symbolic operations.

So in that way he played the same game that I was talking earlier about, taking a phenomena, formalizing it, and enjoying seeing your intuition amplified by the game of symbols.

And then Shannon followed Boole, and he translated...he took Boolean algebra in his master's thesis in 1938, and he just noticed that you can use Boolean algebra to optimize Boolean function on switching circuits for telephone exchanges. That was...for me it was one of the greatest inventions of the 20th century.

You take an abstract symbolic language done by a mathematician in 1850, I believe...oh, no, sorry, 1847, I have Boole's original manuscript so I know 1847, and you apply it to a new phenomena, totally new, relay networks and it started essentially computer science.

I was talking about the same sort of intuition, of seeing one's intuition amplified through formal methods which overrides my entire research. So that's why I would choose Boole and Shannon as my heroes.

But, of course, I would always enjoy Einstein, his humor, and his humanity, and his insights. And Einstein identified the discovery of Galileo, the discovery you can infer cause/effect relationship by experiments to be one of the two greatest achievements of western civilization, or western science, I should say.

The first he attributed to the Greek philosopher for inventing the idea of proof in geometry, and the second was Galileo discovering that with experiment you can verify cause/effect relationships.

I think Einstein is one of my heroes, of course, and I think the field is ready for a third revolution. This is putting the two together, putting experimental science under the auspices of the Greek mathematicians, under the formal guardianship of logic and formal methods.

Stephen: Judea, you continue to make significant historical contributions. How will your growing status contribute to you vision for the world, society, industry, academia, and technology?

Dr. Pearl: This is dangerous, really, because when you have a greater notoriety, or you call it status, people expect you to say something smart and wise on every topic, and I resist it. I should resist, because I do not have anything wise to say on every topic.

For example, I refuse to comment on singularity. At the same time I do have some wisdom in areas that are not usually represented in computer science. For instance, I would try my best to revamp statistical education in this country as well as science education. Science, as I told you before, should be taught as a game between man and nature.

Social science, too. I do have a humble opinion about how the media distorts cause/effect relationships, and I could utilize my expertise in this area. But I generally decline from commenting on a field in which I do not have expertise.

Stephen: This is a bit open-ended. You choose the topic area. What do you see as the top challenges facing us today and how do you propose they be solved.

Dr. Pearl: Education, education and education. That includes science education. It includes social education. For instance, we need to teach pride in the exceptionalism of Western civilization. At the same time, we should teach the kinship of all mankind.

It's very hard, but it's doable, to teach together the dignity of being different and the sanctity of being one. I know it's doable, because my education was that kind, and I think it can be perpetuated to the next generation. I see it being eroded, and I hope it can be restored.

Stephen: It occurs to me that you see human beings, in terms of how we do things, and you've given sort of a representation of how that can be emulated or produced in a robot. Giving the illusion perhaps, of free will, or maybe of consciousness.

Human beings are able to manifest all of this in great leaps in technology, and thinking, and so on. How come we don't see that in other species? How come we didn't see that historically in the Neanderthal?

Dr. Pearl: I didn't get the question. Did see what in the Neanderthal?

Stephen: I would assume though, that humans have this sort of processing of counterfactuals. It's innately ingrained in our machinery, and you've been able to create a representation of this so that it can be also produced in a robot or in a computer.

How come that machinery that we have, and we share much of the genetic code with our ancestor, or with a cousin, the Neanderthal, with a relative again, a chimpanzee, for example. How come we don't get the same kind of outcomes from those other species or predecessors and so on?

Dr. Pearl: I think it was a cultural revolution at a certain point which gave rise to counterfactual thinking in men and societies. I do not have the data. I do not have the knowledge.

I tried to get some linguists to enlighten me in this direction. To tell me whether the history of language shows us that the acquisition of counterfactual vocabulary was accompanied with a certain cultural development. With a quantum jump in the cultural development. For instance, with the development of tools.

I wouldn't be surprised if a linguistic analysis indicates that the two phenomena occurred together. The ability to reason counterfactually, and the development of tools. I'm not sure. I've said I've looked at certain ancient languages, but I couldn't find any analysis on this position.

I only remember that in the Bible, God does not talk counterfactually. Only man talks counterfactually. God asked Adam in the Garden of Eden, "Did you eat from the tree that I told you not to eat from?" He did not ask for "why did you do it?" He asked for the facts.

It's only later. Abraham was the first person in the Bible who talks counterfactually when he comes to God and says, "You are going to destroy the city of Sodom and Gomorrah? Are you going to wipe out the righteous together with the wicked? You can't do that. What if there were 50 righteous people in the city?"

So here you view the first time that the Bible is used counterfactually. What if there were 50 righteous people in the city? Now, this is a dilemma. What is Abraham trying to ask God? Whether he can count properly? He doubts God's ability to distinguish between the wicked and the righteous, or God's ability to count people?

No. I think the idea is that Abraham was after a generic rule. When would a city be destroyed besides the fact that it had x number of righteous people. This generic rule, is a rule of collective punishment, which is applicable not only to the particular city under consideration, but to every other city, to every other community. So, in this way, Abraham was a social scientist who searched for the generic rule behind the phenomenon.

When did science develop? We cannot say that the Neanderthal man was a scientist that had the capability to reason scientifically about phenomena. To seek the generic rule behind behaviors or behind the environment. That was embedded in the language much later. I'm looking for this position.

I don't know the answer. Perhaps this broadcast will reach the ears of some linguist or history of linguistic, or history of language development and will entice their curiosity to look into that question.

Stephen: Perhaps then it's tied to something like the fox 2p gene which enables speech in the way we have it as human beings, or maybe it has to do something with social learning, and the fact that our children learn from us and they can pass that down through generation to generation.

Dr. Pearl: Perhaps. But I think it has to do with credit and blame. You punish your child if your child does something wrong, and it's taken to correct behavior. Dogs can do it too, but not at the linguistic level. See, the child asks, "What have I done wrong?" You tell her, "You spilled the milk on the carpet." And the child says, "I didn't mean to." So next time you become more careful. This kind of conversation doesn't go between monkeys.

Stephen: OK. So again, maybe the root is the fox 2p gene that enables this kind of linguistic manipulation that only humans can do, and that then enables maybe social learning in the context that we're talking about it as well. And that social learning being too counterfactual.

This final question is about your aspirations with the Daniel Pearl Foundation. So perhaps you can give a bit of history, and then answer about your aspirations for that foundation?

Dr. Pearl: Well, the tragedy had driven us to take an action. To utilize all the emotions and good will that the tragedy had evoked, and channel them into one cause, one aim. And this is to roll back the culture of hate, and the ideology of terror that took Danny's life.

And we are trying to muster all the leverages we have in various communities to accomplish that. And we have a great leverage, because the image of Daniel Pearl is immediately recognizable we've learned, all over the world. And when you have such an image that immediately evokes the willingness to act, you can rally people to do good things.

So we are focusing on three areas, this is the journalism, music, and dialogue. We bring journalists from Muslim majority countries to the United States to train them in atmosphere of free press, and to observe the dynamics of American style of journalism, with the hope that they will go back to their countries and generate similar activities.

And we've seen great success for that, and if you take the...if you believe in the saying that investing in one journalist is investing in 100,000 and sometimes a million readers, then you can see the leverage that we have here through the investment in the journalism exchange.

And the other avenue is the music. Danny was a musician, he always carried his violin or mandolin with him in his travels, and he always made friends through music. And he inspired musicians around the world to express themselves, and to use their artistry to mobilize people for this cause.

So we organize hundreds of thousands of concerts [inaudible 01:37:49] during the month of October you'll see in 2,080 concerts. We organized in the [inaudible 01:38:01] in 84 countries, including Iran and Saudi Arabia, and to move people towards understanding of tolerance and humanity.

Concerts may last only one hour, but it leaves a residue if the music is played [inaudible 01:38:22] dedicated to those ideas. It does leave a residue, and we believe that it accumulates. The residue accumulates, it makes people more willing to act in that direction.

And the third avenue is the dialogue. I do take upon myself to...I should say partnership with the Muslim leader, Muslim scholar, with the chair of Islamic studies in the American university, and we travel from city to city and we represent the sentiments of our respective communities.

I will present the Jewish community, and he will present the Muslim community. And we expressed our grievances, and expressed sentiments and try to discover areas of common ground, with the hope that such dialogues will continue in the communities in which we pass.

And they did, and I think we have great satisfaction from showing people that dialogue between faiths are feasible, and can be conducted by two grandfathers in a cordial, friendly way with respect and real understanding [inaudible 01:39:47] avenue.

And of course we have also a series of lectures, and others. But this is the main three tiers through which we're trying to propagate the idea, to roll back the culture of hatred.

Stephen: Judea, you have a lifetime of standing research contributions, influencing diverse domains and with lasting significant global impact, and you've taken the time to share your time, your passion, your history, your success, and all of your experiences through a very, very long globally impacting lifetime of just doing it and making a difference.

And I really do thank you and appreciate you coming in today to do this interview, and to share all of that wisdom with our audience.

Dr. Pearl: Thank you Steve, it's really a privilege for me to share my little knowledge with your audience, and I hope it will be useful.

Stephen: I'm Steven Ibaraki, and this concludes our exclusive interview with Judea Pearl. ACM Turing award recipient in 2012, which is widely thought of as the Nobel Prize of computing. Judea is a legendary pioneer, world renowned distinguished researcher and professor at UCLA.

[music]

Transcription by CastingWords