

Frontiers of Information Technology & Electronic Engineering  
 www.jzus.zju.edu.cn; engineering.cae.cn; www.springerlink.com  
 ISSN 2095-9184 (print); ISSN 2095-9230 (online)  
 E-mail: jzus@zju.edu.cn



## Interview:

# An interview with Dr. Raj Reddy on artificial intelligence

FITEE editorial staff

Zhejiang University Press, Hangzhou 310028, China  
 E-mail: fitee@zju.edu.cn

<https://doi.org/10.1631/FITEE.1700860>

Dr. Raj Reddy is the Moza Bint Nasser University Professor of Computer Science and Robotics in the School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA. He was awarded the ACM Turing Award in 1994.

Dr. Reddy has shown great interest in the ‘New-Generation Artificial Intelligence Development Plan’ issued by the Chinese government in July 2017. In November 2017, while attending a conference in Shanghai, Dr. Reddy accepted our request for an interview in which he gave his views on artificial intelligence (AI). At that conference, Dr. Reddy proposed the idea of cognition amplifiers and guardian angels.

**FITEE editorial staff** (‘Editorial staff’ for short): Comparing the major breakthroughs in AI of the 20<sup>th</sup> century (e.g., a world champion chess machine, mathematical discovery, robotics, speech recognition, expert systems, and computer vision) and the major breakthroughs of the 21<sup>st</sup> century (e.g., language translation, dialogue, autonomous vehicles, deep Q-A, and world champion poker), what is the driving force behind these advances?

**Dr. Reddy:** The early breakthroughs (those of the 20<sup>th</sup> century) in AI were based primarily on knowledge that had accumulated over the last 2000 years and, in particular, the scientific advances of the last 300 years. The breakthroughs that are happening now in the 21<sup>st</sup> century are based on new knowledge that is discovered and used by AI systems themselves. As a result, the new directions in 21<sup>st</sup> century AI are basically the discovery of new

knowledge and its application to solve previously unsolvable problems. The problem with human scientific discovery is that it is a time consuming and laborious process. Early AI systems required human hand-crafting of knowledge in knowledge based systems. Today, machine learning algorithms discover new knowledge (e.g., the hidden patterns) directly from data.

**Editorial staff:** Big data intelligence, cross media intelligence, crowd intelligence, augmented intelligence, and unmanned aerial vehicle intelligence are five technologies featured in AI 2.0. Cognition amplifiers and guardian angels are examples of the use of AI 2.0 technologies. How are cognition amplifiers and guardian angels created in terms of the five technologies featured in AI 2.0?

**Dr. Reddy:** Basically, a cognition amplifier does things that we know how to do—tasks you know how to do but do not want to do because they’re boring. For example, you can drive a car but may prefer a chauffeur or somebody else to drive the car. That is the role of a cognitive personal assistant (i.e., cognition amplifiers). Basically, you should let a machine do anything that is routine or boring. A cognition amplifier needs to collect all the (big) data from different sources. These concepts are similar to both big data intelligence and cross media intelligence in AI 2.0.

Guardian angels do things that you do not know how to do. You do not know that an earthquake is going to happen tomorrow. You do not know that some accident is about to hit you.

In summary, cognition amplifiers amplify what you already know how to do. Guardian angels are agents that do what you do not know how to do or do not have the time and energy to learn, i.e., predicting what is going to happen, events of which you have no knowledge, but that the world can predict.

**Editorial staff:** Learning from scratch is a long-standing ambition of AI research since it can

bypass human capabilities, knowledge which may be too expensive, too unreliable, or simply unavailable. AlphaGo Zero by DeepMind and Libratus by CMU are two examples of systems that learn from scratch. Do you think learning from scratch is a generic approach to artificial general intelligence?

**Dr. Reddy:** My answer is No. There are two aspects of learning from scratch. AlphaGo, poker, and chess do not learn from scratch; they're given a large amount of data. There's another aspect of learning from scratch, which we call incremental learning, where the amount of data is small at first, but grows much larger with predictive learning. Not only that, it can then randomly generate new games that have never been played before to see what happens. In that sense, learning from scratch is not a generic approach. The main difference between AlphaGo and poker Libratus is that when I show you the AlphaGo board (or chess board), all the information that you need to solve that problem is known. When I show you a part of a poker hand, that information is incomplete—there's some other information that is not visible. This happens even in business negotiations, not only in games. I do not know your secret—how much you're willing to pay for my knowledge, business, or product. And there's bargaining going on. I say: "If you want my patent or technology, I want a million dollars." You might wonder: "Would you accept a thousand dollars, or a hundred thousand dollars? At what point will you walk away?" All of that is based on partial information. The missing information may become known only afterwards. It's all completely random chance depending on how much data you are getting and how you process it. All that means is that when AlphaGo, poker, etc. try to examine the search space, in one case all the information is there, and in the other case, there are more cards there in the deck that have not been given to anybody. And those also become part of the search space.

**Editorial staff:** Humans perform reasoning depending on textual, logic, visual, and even intuitive knowledge. As a result, knowledge representation is very important. How do you suggest visual and even intuitive knowledge should be represented?

**Dr. Reddy:** The problem of knowledge representation has been with AI right from the beginning. What is knowledge and how do you represent it?

One way of thinking about it is that knowledge is distillation of information from big data. I begin to see as soon as I am born, I see my mother's face every day. After I have been seeing it for six months, I can recognize my mother's face. But if you appear out of the blue, I cannot recognize your face because I have seen it only once. I have a good representation of the knowledge of my mother's face, but I do not have a good representation and knowledge of your face. It's just a question of a lack of multiple impressions of your face. What is happening is that the human brain is learning the concept of mother's face and different people's faces from visual data over and over again. There is textual data, logical data, speech data, image data, or behavioral data about who is buying what things on Amazon and Alibaba. That data also has to be taken in, analyzed, distilled, and represented.

**Editorial staff:** Interpretable deep learning or explainable deep learning is a new attractive research direction. What do you understand by interpretable AI?

**Dr. Reddy:** That is a question that everybody is asking: how do we explain the output of a deep learning algorithm? What does the data tell you and how did you arrive at this result? In 20<sup>th</sup> century AI, an expert system would write down the rules, 'if this, then that', because that knowledge was known and captured from the head of the experts who were able to explain why that decision was made.

**Editorial staff:** Cross-media intelligence is similar to how a human being thinks. There are many examples of cross-media intelligence, such as video captioning and the visual Turing test. Do you have any suggestions about the future of cross-media intelligence?

**Dr. Reddy:** Cross-media intelligence was first proposed by Prof. Yunhe Pan and has been studied by faculty staff and students at Zhejiang University since 2005. The basic idea of cross media is how to bridge both the semantic gap and heterogeneity gap, and appropriately use the data with multi-modality to understand the surrounding world.

Sometimes, if you do not have additional information from other sources, you are not able to disambiguate the meaning of a word depicted in a picture. So, you need more data, more information from multiple sources (i.e., cross media). In medical

AI, you may take a look at blood test, temperature, height and weight, MRI and X-Ray data. All of these are different sources of knowledge, and you need all that knowledge for a correct diagnosis. So, cross-media knowledge gives you the ability to interpret unambiguously what a person's problem is. That is the main power of cross-media intelligence—it helps you to home in on the problem quickly. It is also called 'data fusion'. In data fusion, different kinds of data are collected; e.g., image data, depth data, and radar imaging data are needed for self-driving cars. A human would use all kinds of things to glean information about their surroundings from their perspective. A camera tells you only one kind of thing. Radar gives you the distance and tells you different kinds of things. When you can see a plane, then you know there's a plane. If you cannot see the plane because of the clouds, radar goes through the cloud, and eventually you can see it. You think there's something where there was nothing before, suddenly for about a second there is an object, and then it disappears.

**Editorial staff:** What is the CMU AI initiative launched in August 2017? What is different about the CMU AI initiative?

**Dr. Reddy:** We do a lot of AI at CMU. There are about 150 faculty members working on many different aspects. The world does not know about them, because we have never brought them together in one place to articulate our big vision for the future. So what we're saying is, at least from a CMU perspective at this point in time, there are at least five or six broad areas related to AI, including machine learning, autonomy, assisting humans, and robotics. We do many things as part of the CMU AI initiative. For AI to reach greater levels of sophistication, experts in each aspect of AI, such as how computers understand the way people talk or how computers can learn and improve with experience, will increasingly need to work in close collaboration. The CMU AI provides a framework for our ongoing

AI research and education.

"AI is no longer something that a lone genius invents in the garage," Dean of the School of Computer Science, Andrew Moore adds, "it requires a team of people, each of whom brings a special expertise or perspective. CMU researchers have always excelled at collaboration across disciplines, and CMU AI will enable all of us to work together in unprecedented ways."



Dr. Raj Reddy is the Moza Bint Nasser University Professor of Computer Science and Robotics in the School of Computer Science at Carnegie Mellon University. During 1960–1963, Dr. Reddy worked as an Applied Science Representative for the IBM Corp. in Australia.

Dr. Reddy's research interests include artificial intelligence and the study of human-computer interaction. His current research interests include: technology in service of society, cognition amplifiers and guardian angels, digital democracy, universal digital archive, voice computing for the 3B semi-literate populations at the bottom of the pyramid, and KG to PG Micro-Universities.

His professional honors include: Fellow of the Institute of Electrical and Electronics Engineers, Fellow of the Acoustical Society of America, Fellow of the American Association for Artificial Intelligence, Member of the National Academy of Engineering, and Member of the American Academy of Arts and Sciences. He is also a member of the Chinese Academy of Engineering, Indian National Science Academy, and Indian National Academy of Engineering. He was the President of the American Association for Artificial Intelligence from 1987 to 1989. In 2011, he was inducted into inaugural IEEE Intelligent Systems' AI's Hall of Fame.

Dr. Reddy was awarded the Legion of Honor by President Mitterrand of France in 1984 and Padma Bhushan by the President of India in 2001. He was awarded the ACM Turing Award in 1994, the Okawa Prize in 2004, the Honda Prize in 2005, and the Vannevar Bush Award in 2006. He served as a co-chair of the President's Information Technology Advisory Committee (PITAC) from 1999 to 2001.