Reflections on Successful Research in Artificial Intelligence: An Interview with Yolanda Gil

Yolanda Gil, Biplav Srivastava, Ching-Hua Chen, Oshani Seneviratne

■ *This article contains the observations of Yolanda Gil, director of knowledge technologies and research professor at the Information Sciences Institute of the University of Southern California, USA, and president of Association for Advancement of Artificial Intelligence who was recently interviewed about the factors that could influence successful AI research. The editorial team interviewers included members of the special track editorial team from IBM Research (Biplav Srivastava, Ching-Hua Chen) and Rensselaer Polytechnic Institute (Oshani Seneviratne).*

Editorial Team: What are your observations on how has the nature of AI research changed and evolved over the years?

Gil: I worked on my thesis in the 1980s, when AI research methodology was very different than it is today. Back then, AI research was very broad and cross-disciplinary. The machine learning conferences I attended would include cognitive psychologists developing new capabilities for education and learning, human computer interaction researchers doing work on knowledge acquisition, and mathematicians focused on the statistics of machine learning and formal aspects of learning. Because of this, AI researchers like myself would naturally mingle with researchers from other disciplines. In contrast, today we see research that is more narrow and focused. This could be a sign of a maturing field, because sometimes when a field grows people tend to specialize into very focused areas. But having a narrow view of AI is not necessarily a good thing in my opinion.

 ISSN 0738-4602

Editorial Team: How does one evaluate AI research?

Gil: For me, the evaluation of research in AI depends on the disciplinary angle that is taken. There are significant differences in evaluation and in methodology for AI research that focus on human computer interaction, or cognitive science, or physics, or social networking, or philosophy. For example, a cognitive psychology methodology may place a lot of emphasis on theoretical models of the mind and implementing AI systems that approximate them well based on particular measurements from user studies. In contrast, a more mathematical methodology may evaluate accuracy based on some statistical metric and may be more focused on algorithmic improvement. So different evaluation methods are appropriate for different kinds of AI research.

Editorial Team: Is it difficult to learn how to understand and appreciate different disciplinary angles?

Gil: You have to learn enough about the other discipline, which can be a challenge. I'll let you in on a little secret: Long plane rides are a great opportunity to learn something new and different. Whenever I am on one, I really cherish the opportunity to read without interruption. Before the trip, I will prepare a big folder of materials and books to teach myself about a topic that I am less familiar with. Reading them makes the trip feel very short and I can learn something about a new discipline or topic that I have been curious about!

Editorial Team: What factors do you see inhibiting the rate of success in AI research?

Gil: We need to emphasize more the importance of reproducibility, which is a fundamental aspect of the scientific method. One key component of reproducibility is sharing and publishing (with DOIs) all of your data and your code together with execution details (such as parameters and key intermediate results). This is important because it helps document the specific work that was done. It helps other researchers understand the context of the claims you are making. It also helps others build on what you did and vice versa. This is something I have been promoting and pushing in the scientific-paper-of-the future initiative.[1] This initiative promotes a future in which all relevant details of a computational experiment are exposed properly and in a structured way.

A positive trend that we are seeing now is that more and more people want to share their results in the open as soon as possible. The cycle of waiting for six months to see your paper come out is too long. People are publishing their work on *ArXiv* so that their results can be out and then they move on to the next idea. There is an incredible eagerness by young researchers to disseminate their work through public sites. We need to encourage this and do it for data and software as well, for every research product in addition to papers.

Editorial Team: Many countries have come up with a national AI policy. What are your thoughts on government involvement in AI research and development?

Gil: The government can play an incredibly important role. Through its funding programs it can sustain basic research. The Internet came out of government investment, and so did technologies like chat bots and search engines, all of which are pervasive in our society today. So, we are reaping the benefits of investments in basic research that were made by the government decades ago. Sometimes government funding agencies, perhaps influenced by the interests of the commercial sector in funding the next start-up, seem very focused on projects that have immediate applications. I think it's really important for the government to retain a focus on long-term problems and to consider sustained decadal-scale funding programs for the most challenging AI topics.

Editorial Team: What skills are essential for young AI researchers to have?

Gil: I teach my students to be very mindful of long-term, deep ideas. Students will often focus on the next paper, or on the things they can do today to extend some code they are working on. However, one of the most important skills for any scientist to have is to think deeply about a very challenging problem. To do that, you need to invest the time. You need to make a commitment to spend time defining and thinking about that problem. It is much harder to think about what will be important in 10 years than it is to think about the next paper or the next line of code. It is very hard to devote yourself to thinking rather than coding, but it is very important if we want fundamental advances in AI. For students, this is especially hard, and as professors, we need to teach it.

This is also why it is important for students to go to a broad conference like Association for the Advancement of Artificial Intelligence (AAAI), to get a good view of different research topics in AI, which can promote deeper thinking about problems. For example, at the AAAI conference, you will find sessions on machine learning and natural language processing, others on robotics and human computer interaction, and others on reasoning and planning. And many of the papers will combine research on several areas of AI. You will find both interesting applications and fundamental research. I also attend specialized conferences, but I find the AAAI conference most inspiring for formulating long-term research problems.

Finally, a problem I find that it is when I am I thinking about long-term topics that I am most excited, work the hardest, and get the most satisfaction at a personal level.

Editorial Team: Because deeper ideas take time to develop, and publications may not come about so quickly, what are some methods you use to know that you are making progress against deeper ideas?

Gil: This is how I measure progress: I start a folder on a far-fetched topic and watch how it grows over time. Initially the folder is very thin, but I keep adding to it as I think more about the topic. Six months later it has expanded with details, and in yet another 6 months it will have expanded further, and so on. At some point, I will have enough of an idea to work with a student, and eventually, I'll understand the problem enough to formulate a 5-year

plan. This is a many-month process, and on day 1 you don't see how you could possibly measure improvements, but eventually something clicks, and you can really see the path forward.

Editorial Team: Many companies are now interested in commercializing AI and creating their own narratives around AI. What are your thoughts about what you see?

Gil: I have been doing AI research for more than 30 years, and I have always seen companies presenting AI in a way that promotes their commercial interests. This is natural and is nothing new. In the 1980s, all the banks wanted investors and customers to know that they were using expert systems and they had AI capabilities in-house. As an AI researcher, I like to take the opportunity to reflect on whether I agree with a particular company's approach to AI. For example, do I like the ethics they use when they release AI capabilities? As AI researchers, we have tremendous responsibility to work for (or with) companies that we are in agreement with and whose narratives for AI we find compelling.

Editorial Team: As the incoming president of AAAI, what influence do you hope to have?

Gil: I view a position like this in an organization like AAAI as a position of service. One is there to serve the community, to really tap into what the community wants and needs. I may have some ideas but I am very much interested in understanding how the community wants to move forward. One aspect that the AI community cares about deeply is the connections between industry and academia. We have the Innovative Applications in AI conference during AAAI, and that is very exciting but we need additional efforts in this area. For example, AAAI is one of the founding board members of the Partnership for AI, which brings together companies with academic and scientific organizations to understand important issues concerning the practical uses of AI. Another topic that our community really cares about is education. We have the Educational Advances in AI conference during the AAAI conference, which brings together professors and teachers in a wide range of institutions for higher education. AAAI just started a new initiative on K-12 education, which is incredibly important. Not only do we need to build a pipeline of researchers in AI, but we have to recognize that a K-12 student today will be a consumer of AI tomorrow, and a voter (possibly even a member of Congress!). We want all students to understand the big questions in AI and provide them with the ability to think critically about how AI systems work and how they may affect their lives. Diversity and inclusiveness in AI is incredibly important and a very effective way to address this by making AI accessible to all students in early education.

Editorial Team: What types of AI research should we do less of? Or more of?

Gil: I think AI researchers are extremely creative and diversified in their topics of research, and this is very healthy for our field. So, let me talk about an area where I think we should do more: Our AI systems have no notion of their own limitations. They don't know what their capabilities are or what exists beyond their capability. In the human world, if you walk into a pharmacy and ask someone for a flu shot, if that person is a cashier or a receptionist, they will tell you why they can't give you one and they will direct you on who to go to instead. Our AI systems are rarely designed to do this. When asked a question, our AI systems will give you an answer, but they don't know the context, and they often don't know the risks involved. I wrote a paper recently about thoughtful AI[2] that talks about how an AI system could have more awareness of how it fits into the world so that it could say "I don't know a lot about that so I cannot be helpful, but here's how you can find help." I wish there was a lot more research on this topic.

Editorial Team: What excites you most about your current research?

Gil: One of the things we are able to do much better today than 5 years ago is to capture information about a scientist's data analysis process and represent it as a semantic workflow. This year we have been studying how data analyses are performed in hydrology and agriculture modeling to manage water resources and food production. As we continue to study different scientific disciplines, we see more commonalities. Semantic workflows will allow us to teach machines to help us make scientific discoveries. And, I don't mean giving a machine data and asking it to come up with correlations. I mean teaching a machine to actually design an entire approach to test hypotheses by finding appropriate data, performing data analysis, and making decisions about what results may be significant. This is how a scientist would approach a problem. Today, AI systems are not full-fledged scientists, but they are beginning to be able to design approaches to data analysis. I believe we are on the brink of seeing AI systems deeply transform how we approach scientific discoveries.

Editorial Team: Thank you for sharing your perspectives with us today, Dr. Gil.

Gil: You're welcome, I really enjoyed your thoughtful questions.

Notes

1. www.scientificpaperofthefuture.org/.
2. The paper is available at doi.org/10.3233/DS-170011.

Yolanda Gil is director of knowledge technologies and research professor at the Information Sciences Institute of the University of Southern California, USA, and president of AAAI.

Biplav Srivastava is a distinguished data scientist and master inventor at IBM's Chief Analytics Office.

Ching-Hua Chen is a research staff member at the IBM T.J. Watson Research Center in Yorktown Heights, New York.

Oshani Seneviratne is the director of health data research at the Institute for Data Exploration and Applications at the Rensselaer Polytechnic Institute.