The Web Conference Keynote*
AI Grand Challenges: Past, Present and Future

Ganesh Mani
Carnegie Mellon University
ganeshm@andrew.cmu.edu

Abstract

Everybody has opinions on grand challenges—bold tasks that capture the imagination of researchers and system builders—for AI. One of the earliest set of challenge goals was enunciated by Turing award winner and AAAI President, Raj Reddy, in 1988. I attempt to provide an accounting of the progress that has been made in the field, over the last three decades, towards those challenge goals. While some tasks such as the world-champion chess machine were accomplished in short order, many others remain incomplete. A new set of challenges for the current decade are also proposed, spanning the Health, Wealth and Wisdom spheres. Bridging the gap from narrow AI to general-purpose AI will be required to solve some of these challenges; teaming architectures in the service of humankind will also need to be emphasized. The mantra should be: Of the people, by the people *with machines*; for the people!

1 Reddy's Challenges in 1988

The challenge tasks laid out by Reddy [1988] touched upon everyday elements—spanning communication, transportation and games—plus infrastructure requirements (on earth as well as for space explorations).

The world champion chess machine challenge was accomplished in short order while the journey is only half-complete as far as the accident-avoiding (self-driving) car goes. The translating telephone task can arguably be deemed complete as smart phones with translation apps—that are able to handle dozens of languages—have proliferated (via technology platform companies such as Google, Microsoft and Facebook). AI has aided minor mathematical discoveries but a major result that is hitherto unknown to human mathematicians is still elusive.

In the realm of self-organizing systems, the original goal called for acquiring significant capabilities via perception-mediated learning and discovery. Learning from textbooks like youngsters do or assembling an appliance by watching a human mechanic. The *Aristo* project from the Allen Institute for AI [Clark et al., 2020] reports a 90% performance metric on the NY Regents 8th Grade Science Exam. While the vocabulary comprehended is significant, we are still in the realm

*I am grateful to Marko Grobelnik for the invitation to speak.

of non-diagram, multiple choice questions. Developments such as IBM Watson's winning performance on *Jeopardy!*, the family of transformer models starting with BERT and GPT-3's text generation abilities while impressive still make many people in the field wonder about the level of true understanding and generality (especially as it relates to communicating with a human; or being robust to facts implied by the text as well as different types of re-wording of the queries or questions). Moderate amounts of progress is how I score the overall challenge.

The self-replicating systems challenge was foreseen as a way to perhaps transport "seeding" robots and a minimal set of tools and then to leverage in situ resource utilization (e.g., on Mars or on an island or region cut off due to a natural disaster). A modicum of progress has been made and performance teams are working on various aspects, for example, as part of the NASA Space Robotics Challenge.

The World Wide Web has permitted facile indexing and fast retrieval for widespread sharing of information. News organizations post digital content in real-time and social media platforms permit the rapid addition of user-generated content. This poses new challenges: how to discern the veracity and source authority of a news story, separating facts from opinions, summarizing news stories, and highlighting any unique details a particular news article may provide. Reddy has opined and I concur that the ideal information transmission goal—of getting the right information at the right time, in the right granularity and language to the right person—is still a work in progress.

2 New Challenges, Commentaries and Feedback

For the current decade, I introduce a new set of potential challenges spanning the Health, Wealth and Wisdom spheres; progress towards them will require technical accomplishments as well as deliberations around policy implications and societal impact. I also invited half a dozen thought leaders with varying vantage points—involved in different aspects of AI, including influencing funding towards the field—to opine and suggest Grand Challenges; their commentaries are featured in the sidebars of our article [Mani et al., 2021]. The article includes a number of references—in its endnotes and citations—which reflect the state of the art in the field and other nuanced perspectives around challenge tasks. Human-machine teaming and ethics—topics growing in importance in the field—are also addressed therein.

The challenges proposed for the current decade include winning the *New Yorker* Cartoon Caption Contest (multiple times and with explanation); and, for an AI system to try and convince humans that it is trustworthy, as its decisions are fair and ethical.

Here, I will briefly add to the discussion to dovetail with the oral conference presentation[1] and also describe some of the reactions our article sparked.

Selman et al. [1996], during a panel discussion at the AAAI-96 conference, also laid out a set of challenges that they termed "ambitious and interesting, yet not so open-ended." Many of those goals continue to be germane, remaining open for completion. Deep learning expert and recent Turing award winner Yoshua Bengio (in a personal communication) suggested tackling climate change and pathogen resistance as significant challenges. He also enumerated some method goals that would be worthwhile for the field to focus on: out-of-distribution generalization, learning to

[1] https://www.youtube.com/watch?v=FkH3T7guZ6Y&t=6060s

reason causally and discover unobserved latent causal concepts; as well as knowledge acquisition in a complex high-dimensional environment with an exploration policy.

Government funding agencies and academies in the US—for instance, the NSF[2], the DOE's AITO Office[3] and the NAE[4]—have also proposed their own versions of challenge tasks and goals.

Arguably, the most thought-provoking feedback originated from a physician acquaintance—a power user of AI-assisted radiology—who asked how soon could AI significantly help individuals with activity limitations, as he has long been caring for a son with cerebral palsy.

3 United Nations Sustainable Development Goals

Since the keynote coincided with Earth Day 2021, I thought it was important to address sustainability guidance for human-machine teams. A candidate lodestar is the set of Sustainable Development Goals (SDGs) recommended by the United Nations. The 17 goals include "no poverty," "reduced inequalities," "climate action" as well as "peace, justice and strong institutions." It can be thought of as a call to action to protect the planet, and to ensure peace, harmony and prosperity for all by 2030. Congruent with the above, we become smarter by harnessing AI and other advanced technologies especially towards these broader thematic priorities: one, doing no harm, by taking steps to preempt all kinds of warfare, be it kinetic, cyber or biological; two, preserving democratic institutions and the rule of law; three, providing contemporary, free K-12 education for all, so that where a child is born does not dictate how far she would go in life and how long she would live; four, exploring geoengineering schemes—for instance, afforestation or reforestation, solar radiation management, as well as carbon dioxide and other greenhouse gas removal—to turn the clock back on climate change.

UN Secretary-General Antonio Guterres, in his Earth Day message, called for bold climate action to limit global temperature rise to no more than 1.5 degrees Celsius. However, AI may not be an instant panacea for all the challenge use cases above. Vinuesa et al. [2020], using an expert elicitation process, have concluded that AI is conducive to supporting 134 targets spanning all the SDGs, while potentially having a negative impact on 59 of the targets. Employing AI to nudge human behavior towards sustainability— suggesting greener alternatives as individuals and institutions make consumption or spending plans and decisions—should be another research focus.

4 Concluding Thoughts

Prizes (along with bragging rights) can motivate researchers and innovators to make progress towards solutions to hard challenge problems, spurring innovation. Francesca Rossi also astutely noted in [Mani et al., 2021] that grand challenges should not be considered in isolation, as working on one will bring new insights into many other ones.

AI projects are best thought of as a triad: data (including alternative data), domain knowledge and tools. As open-source libraries underlying the tools become more widely available, the inno-

[2] www.nsf.gov/news/special_reports/big_ideas

[3] www.energy.gov/articles/secretary-perry-stands-office-artificial-intelligence-and-technology

[4] www.engineeringchallenges.org

vation locus is shifting to data choices and know-how (including the selection of tools that match the data characteristics and modeling goals). In addition, how multiple subsystems - that may be specialized - are architected and assembled may be key to the elusive goal of general-purpose AI. All indications are that it will not be adequate to simply string together a bunch of today's narrow-AI systems, that work very well on a few standardized datasets.

For instance, the bulk of everyday tasks that a 10-year old child performs—many of them seemingly effortlessly—and the reasoning associated with them is out of reach of today's AI systems. While deep-learning components may be suited for perception and reflexive motor reactions (much like a very young child withdrawing a finger from a hot surface), a motley collection of building blocks may be required to accomplish even simple tasks that exhibit commonsense (e.g., noticing that an electric stove is not plugged in and hence concluding that it is likely not hot). The higher-level building blocks will most probably rely on perceptual grounding of language and symbols, what-if simulations via digital twins of humans and causal reasoning that encodes many aspects of commonsense. Enabling facile communication among humans and machines in a *lingua franca* is also key. In the short term, multi-human, multi-machine teams will likely exhibit the strongest competence—as well as prudent risk management, when encountering *IROPs*—in many milieus. Other than AI, quantum computing and behavioral science advancements may further enhance the capabilities of such teams.

References

Peter Clark et al. From 'F' to 'A' on the N.Y. Regents Science Exams: An Overview of the Aristo Project. *AI Magazine*, 41(4):39–53, December 2020. doi: 10.1609/aimag.v41i4.5304. URL https://doi.org/10.1609/aimag.v41i4.5304.

Ganesh Mani, Frank Chen, Steve Cross, Thomas Kalil, Vanathi Gopalakrishnan, Francesca Rossi, and Ken Stanley. Artificial Intelligence's Grand Challenges: Past, Present, and Future. *AI Magazine*, 42(1):61–75, Apr. 2021. URL https://ojs.aaai.org/index.php/aimagazine/article/view/7375.

Raj Reddy. Foundations and Grand Challenges of Artificial Intelligence: AAAI Presidential Address. *AI Magazine*, 9(4):9, Dec. 1988. doi: 10.1609/aimag.v9i4.950. URL https://ojs.aaai.org/index.php/aimagazine/article/view/950.

Bart Selman, Rodney A. Brooks, Thomas Dean, Eric Horvitz, Tom M. Mitchell, and Nils J. Nilsson. Challenge Problems for Artificial Intelligence. In *Proceedings of AAAI-96, Thirteenth National Conference on Artificial Intelligence, Portland, Oregon, August 1996. AAAI Press, Menlo Park, California*, pages 1340–1345, October 1996. URL https://www.microsoft.com/en-us/research/publication/challenge-problems-articial-intelligence/.

Ricardo Vinuesa, Hossein Azizpour, Iolanda Leite, Madeline Balaam, Virginia Dignum, Sami Domisch, Anna Felländer, Simone Daniela Langhans, Max Tegmark, and Francesco Fuso Nerini. The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11(1), January 2020. doi: 10.1038/s41467-019-14108-y. URL https://doi.org/10.1038/s41467-019-14108-y.