AI RFI Responses, October 26, 2018

_____

Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan RFI Responses

Computing Community Consortium (CCC) Response to NITRD
"RFI on Update to the National Artificial Intelligence Research and
Development Strategic Plan"
Elizabeth Bradley, Ann W. Drobnis, Gregory D. Hager, Mark D. Hill, David Parkes

This response was prepared by the Computing Community Consortium (CCC). The mission of the CCC is to catalyze the computing research community and enable the pursuit of innovative, high-impact research. Our goal is to identify and call attention to major research opportunities for the computing community.

The National Artificial Intelligence R&D Strategic Plan outlines the continued need for basic research in Artificial Intelligence (AI), a focus that is crucial for our continued leadership amid increased international investment and competition. However, the current plan is two years old, and this is a field that is moving very quickly. We are pleased that the NITRD NCO and the Select Committee intend to update the Strategic Plan to reflect current priorities. The CCC concurs with the assertions that advances in AI will have transformational impacts on society, and that sustained support for fundamental research in a wide array of AI-related topics will continue to be crucial in enabling those advances. To that end, and of relevance to your work, the CCC has embarked on the development of an AI Research Roadmap, with broad input from the computing research community, with the goal of identifying key research areas in AI and delivering its findings by Spring 2019. We hope that the NITRD NCO and the Select Committee will view this roadmapping effort as an input from the computing research community, represented by CCC, to the updated Strategic Plan.

Recent advances in machine learning may create the impression that the field is rapidly shifting from fundamental work to application—as evidenced by a rapidly unfolding landscape of new applications and fueled by unprecedented commercial investment. However, fundamental questions are very challenging and require sustained long-term investments that are not happening, so the next wave of transformative ideas in AI will be very slow coming. Furthermore, the supporting science underlying truly reliable, replicable, and scalable AI-based systems is still very much needed. This foundational research should be informed by the current wave of empirical progress, but will be even more essential as current progress slows and the next generation of AI paradigms are needed.

There have been several historical waves of excitement and progress in AI, each driven by a new family of technologies. For example, expert systems in the 1980s became popular as a mechanism to capture knowledge. Likewise, advances in mapping and planning technologies in the 1990s created excitement in robotics. Machine learning began to drive advances in speech,

natural language processing, and vision in the early 2000s, which has been accelerated by ever more capable systems built around deep learning.

The long-term promise of AI is to produce systems that understand and interact with their environment in intelligent ways. That system may be an assistive robot for independent elderly living, a personalized math tutor, an automated driving system, a healthcare decision support system, or a flexible robot on a manufacturing floor. While learning technologies will likely contribute to many of these applications, they constitute only part of a truly intelligent system.

When considering this intelligent system perspective more deeply, there are a few general attributes that stand out:

1. An intelligent system will likely involve many components with different functions. Very few systems involve a single AI technology; integration of AI functions with both AI and non-AI elements is necessary to create almost any real-world system.
2. AI systems will be embedded in a context or an environment that often includes people and other AI systems. The definition of what constitutes success is heavily informed by that context—e.g. how reliable the AI must be, how well it can perceive and understand the context, how adaptive it should be, what types of interfaces it needs, what types of actions are possible and appropriate, what behavior is possible with incomplete knowledge and under uncertainty, and so forth. There is no generic definition of "correct" or "reliable."
3. The scope of pre-programmed knowledge vs. learning will vary widely. Understanding complex signals such as speech or vision is fundamentally learning-driven today, but organizing knowledge and applying it in reasoning processes is not. In some cases, physical models (e.g., of car dynamics) are well understood and should be applied; in others (how a soft pillow behaves) are not and may need to be learned.
4. Continued AI advances, especially involving substantial data, will require concomitant advances in computer systems for faster and less expensive cycles, communications, and storage.
5. Repositories of shared data and shared knowledge about the world, curated with appropriate privacy regulations, and made accessible to the research community are critical for progress, especially in areas where obtaining data is complex, such as with vulnerable user populations (children, elderly, individuals with differences and special needs).
6. AI and ML systems could benefit from inclusion of powerful interactive visualization tools to detect faulty and biased data, understand the program operation to remove bugs and improve performance, monitor outcomes so as to ensure safety, and support explainability and continuous improvement.

7. Ethics, privacy, and security are complete research areas on their own right, but for AI systems of the future to not only be adopted, but to be fully utilized, we must consider the importance of these topics to AI research and development.

There is no single "silver bullet" that will drive AI. Machine learning comprises a very important set of methods; one of those, deep learning, has been successfully applied to a particular type of problems. However, deep learning is limited, and cannot be applied to many more general problems of reasoning in AI.

Strategy 1 of the NITRD plan clearly lays out the need for continued long-term investments in AI research by the Federal Government. We cannot stress how important this is; even though many companies are doing work in AI, they are doing work that is in line with their particular needs. It is only through continued basic research that we will see transformational breakthroughs in AI for high-risk, high-reward situations.

There is much concern—particularly outside the research community—around AI replacing humans, and the effect that may have on the labor market. However, we feel strongly that there are great opportunities for AI to *augment*, not *replace*, humans as suggested by Strategy 2 in the Strategic Plan. For example, a recent panel[1] at the American Association for the Advancement of Sciences (AAAS) in 2018 outlined many areas where augmentation with AI techniques would have a substantial economic or social benefit, including healthcare, transportation, manufacturing, and a variety of service industries. To create technologies that can address these sectors, fundamental advances are needed in learning, planning, human-machine interaction, perception, multi-agent systems, and speech, just to name a few key AI research areas.

While much work is being done in AI, both in research and application, by academia and industry, the CCC believes that this complex collection of issues demands a concerted effort in establishing a set of research priorities and long-term investments. As part of that effort, the CCC is working with leading researchers in the computing research community to gather academic and industrial researchers across disciplines to produce a roadmap[2], culminating with a report in the Spring of 2019. This AI Roadmap Effort is led by Yolanda Gil (USC) and Bart Selman (Cornell). In preparing for this effort, the CCC and Roadmap Organizers have used the current Strategic Plan as a starting document. This effort involves a range of activities, ranging from listening sessions at major conferences to focused workshops:

- The first workshop is on **Integrated Intelligence**, led by Marie desJardins (Simmons) and Ken Forbus (Northwestern). In this workshop, areas to be explored include

---

[1] https://cra.org/ccc/ccc-aaas-2018/#ai
[2] https://cra.org/ccc/visioning/visioning-activities/2018-activities/artificial-intelligence-roadmap/

understanding the mind, composing intelligent capabilities, and open repositories of world knowledge.

- The second workshop is on **Interaction**, led by Kathy McKeown (Columbia) and Dan Weld (U Washington). In this workshop, areas to be explored include interactions that matter, trust and responsibility, and people interacting online.
- The third workshop is **Learning and Robotics**, led by Tom Dietterich (Oregon State) and Fei-Fei Li (Stanford). The topics to be explored include deeper learning, integrated statistical learning and symbolic representation, and diversified learning modalities.

As a community organization working on behalf of the computing research community for over 12 years, the CCC has seen firsthand how important Federal Strategic Plans can be in helping to create technologies that impact national priorities and societal needs. However, we have also seen great plans written with no action. The Administration has taken an important first step in acknowledging the rapid pace of advancement in the field and directing an update of the AI R&D Strategic Plan. We strongly suggest that not only should this plan be periodically updated, but that a periodic review of its effectiveness should be considered. This would seem a natural fit for an organization like the President's Council of Advisors for Science and Technology (PCAST) – or other entity charged with taking on the responsibilities of the President's Information Technology Advisory Committee (PITAC) – just as they have overseen the Federal investment in IT R&D through a biennial review of the NITRD program. We believe this would be a reasonable extension of Strategy 6 in the Strategic Plan.

These are our intermediary recommendations and suggestions for the planned update, but we hope that the NITRD NCO and Select Committee understand that these suggestions are a snapshot at this point in time, and not reflective of the full community voice, which will be more fully represented in the report that follows from the community workshops and listening sessions over the next few months.

We appreciate this opportunity to provide input into this important update the nation's AI R&D Strategic Plan. We intend to keep the NITRD NCO and the Select Committee apprised of our AI research roadmapping effort and hope it will provide additional useful guidance as you work to set the course for Federal investments in and oversight of AI R&D.