Computing Community Consortium
OSTP AI RFI Response
July 2016

Overview

On June 7, 2016, the Computing Community Consortium (CCC) and Association for the Advancement of Artificial Intelligence (AAAI) hosted a roundtable discussion on Artificial Intelligence for Social Good. The following response is a summary of the roundtable discussions. A more thorough report will be published by the CCC later this summer, and available at http://cra.org/ccc/resources/ccc-led-whitepapers/. The remainder of this document is organized into a brief recitation of discussions for the four areas discussed in the workshop, followed by some cross-cutting observations and recommendations.

AI for Urban Environments

The urban computing workshop session focused primarily on transportation networks, the goal being to use AI technology to improve mobility. However, transportation can be viewed as a concrete example of a service industry which is very likely to be transformed over the coming decade by AI technologies. Time spent commuting to school or to work is time not spent studying or with one's family. Lack of transportation reduces access to preventative healthcare, easy access to supermarkets with healthful food is highly correlated with obesity (and hence heart disease, diabetes, etc.) and easy access for people to standard bank accounts is costly. AI technology has the potential to significantly improve mobility, and hence substantially reduce these and other inefficiencies in the market. Technology exists that can mobilize people who have been immobile; to increase flow/decrease congestion; and autonomous vehicles have the potential to decrease emissions. The easier it becomes for people to move about, the more vibrant our urban areas will be; likewise, the more fruitful the social and economic interactions that take place inside them will be.

Ubiquitous connectivity and instrumentation are enabling us to measure things that were previously unmeasurable. We can now collect information about individuals' travel patterns, so that we can better understand how people move through cities, thereby improving our understanding of city life. AI technology can then be leveraged to move from descriptive models (data analytics) to predictive ones (machine learning) to prescriptive decisions (optimization, game theory, and mechanism design). The potential of this transformation is being demonstrated in pilot systems that optimize the flow of traffic through cities, and in new on-demand, multi-modal transportation systems. It is now within the realm of AI technology to optimize traffic lights in real time, continuously adapting their behavior based on current traffic patterns (Smith, 2016); and to dispatch fleets of small vehicles to provide on-demand transportation, address the "first and last mile" problem that plagues many urban transit systems (Van Hentenryck, 2016). More pilot deployments are needed to fully understand the scope of the transformation that is under way in our cities.

In spite of the significant promise, many challenges lie ahead before these new opportunities can be fully realized. Transportation systems are complex, socio-technical systems that operate over multiple spatial and temporal scales. It is critical that we scale up existing pilots to multi-modal transportation models -- incorporating pedestrians, bicycles, cars, vans, and buses -- so that we can begin to understand how these models will impact big cities. Fundamental to this effort, it is crucial that we understand the human behavioral changes that new forms of mobility will induce, and the impact those behaviors will have on the efficacy of our system.

Sustainability

Sustainability can be interpreted narrowly as the conservation of endangered species and the sustainable management of ecosystems. It can also be interpreted broadly to include all aspects of sustainable biological, economic, and social systems that support human wellbeing. In this panel, the discussion focused primarily on the ecological component, but the larger issues of social and economic sustainability must be considered as well. Automated data collection systems develop and deploy sensor networks (e.g. Trans-Africa Hyrdo-Meteorological Observatory; www.tahmo.org), camera traps to collect image or acoustic data, or unmanned aerial vehicles to obtain video imagery. AI algorithms are applied to optimize the locations of these sensors and traps. Crowd- sourcing and/or employing technically-trained people to collect data, such as the freshwater stream surveys conducted by the EMAP project (https://archive.epa.gov/emap/archive-emap/web/html/), are being married with computer vision methods as another hybrid method of data collection.

Techniques from data mining, statistics, and machine learning are used to discover trends and fit models. Such models can predict migration, dispersal, reproduction, and mortality of species. Virtually every ecosystem management problem combines an ecological model with an economic model of the economic costs and benefits of various policy outcomes. Examples include the design of a schedule for purchasing habitat parcels to support the spatial expansion of the Red Cockaded Woodpecker (Sheldon, et al., 2010; Sheldon, et al., 2015), and the use of detailed bird migration models developed by the Cornell Lab of Ornithology to rent rice fields in California (Nicol, et al., 2015). Algorithms for computing these policies combine ideas from network cascade analysis (maximizing spread in social networks) with techniques from machine learning, AI planning and decision-making, and Monte Carlo optimization. Finally, the PAWS project (Fang, et al., 2016) applies AI algorithms for game theory to optimize the patrol routes of game wardens in order to maximize their deterrent effect while minimizing costs.

A major challenge for the medium term is to develop methods that can collect and model data encompassing a broad range of species at continental scales. A related challenge for current modeling efforts is that they generally assume stationary (steady-state) climate, land use, and species behavior whereas the real systems are experiencing climate change, rapid economic development, and continuing evolution, dispersal, and natural selection of species. Furthermore, as the scale of policy questions grows, it is no longer possible to focus only on the biological components of a system. Instead, one must incorporate models of social, cultural, and economic activity. Finally, sustainability hot spots are often located in developing countries. Issues that arise include poor networking infrastructure, little access to high-performance computing resources, lack of local personnel with sufficient education and training, and persistent corruption.

In the longer term, we must confront the fact that the long term behavior of ecological, economic, and social systems is radically uncertain. How can artificial intelligence methods deal with the uncertainty of these "unknown unknowns"? When formulating and optimizing management policies, we should adopt risk-sensitive methods. This is an active area of research (see, e.g., Chow, et al., 2015), and much more work is needed to understand how we can ensure that our models are robust to both the known unknowns (as in traditional risk management methods) and the unknown unknowns.

Healthcare

AI is well-positioned to have a broad and sustained impact on many aspects of healthcare. Social media analytics is emerging as an alternative or complementary approach for instantly measuring public health at large scale and with little or no cost. The nEmesis system, for example, helps health departments identify restaurants that are the source of food-borne illness (Sadilek et al. 2016). Decision support in a clinical environment is a second important area. The Surgical Critical Care Initiative (SC2i), a Department of Defense funded research program, has deployed two clinical decision support tools (CDSTs) to realize the promise of precision medicine for critical care (Belard et al. 2016). The invasive fungal infection CDST was deployed in 2014 to assist military providers with treatment decisions both near point of injury and at definitive treatment centers. The massive-transfusion protocol (MTP) CDST is currently being assessed under a two-year clinical trial at Emory-Grady, one of the two SC2i civilian hospitals. Automated real-time surveillance tools, operating from the electronic health record, identify individuals at risk for severe sepsis and septic shock at the early stages of decline, and much earlier than standard of care (Henry et al., 2015).

Opportunities in this space include:

Targeted therapy decisions: Many chronic diseases are difficult to treat because of high variation among affected individuals. Computational subtyping, for example, seeks to refine disease definition by identifying groups of individuals that manifest a disease similarly (Saria & Goldenberg, 2015) (Collins, 2015). These subtypes can be used within a probabilistic framework to obtain individualized estimates of a patient's future disease course (Schulam & Saria, 2015).

New sensors, new healthcare delivery: AI can be used to analyze social media data and discover and suggest behavioral and environmental impacts on health -- e.g tracking influenza or quantifying alcohol and drug abuse in communities. Social networks can also be used to address the informational and psychosocial needs of individuals and the opportunity for cost-effective interventions for addressing mental health, addiction, and behavioral health issues using modern low cost sensing technologies. Low fidelity sensors, some of which are diagnostic, together with AI and internet technologies can enable low barrier telemedicine for example for chronic healthcare. Advances in natural language processing and machine reading can be used to synthesize, integrate and appropriately disseminate new medical knowledge (e.g., as reported in journal articles).

Pivoting from personalized medicine to personalized health will keep people from going to the hospital in the first place, and dealing with life issues and not just specific disease. For this, we need to move to modeling of the health of individuals and populations by using integrated data sets--- electronic health records data and other data gathered within the health system with genomic, socio-economic, demographic, environmental, social network and social media and other, non-traditional data sources, such as social service and law enforcement data.

Collaborative Decision-Making approaches that allow decision makers to reason with models of the health of individuals are needed. For example, can a healthcare provider ask how would a health trajectory change if the individual was being treated with two different drugs?

Challenges include: 1) addressing Bias that arises in fitting models from observational health data sources; 2) privacy and security methods that support work with data in a way that both sustains its utility while the decisions and outcomes of working with the data do not reveal information about individuals is essential; 3) incentive alignment to encourage various actors in the health ecosystem a reason to collect additional data and make their data available to the rest of the healthcare ecosystem; and 4) cloud-based data science platforms and common data models should be developed and promoted in order to reduce the barrier to entry for researchers and increase the likelihood of societally beneficial outcomes.

Public Welfare

AI has not had a lot of impact on fundamental issues our society faces today. However many opportunities exist. For example, the University of Chicago partnered with Chicago Department of Public Health to build a system to predict which children are at risk of lead poisoning to allow CDPH to deploy inspectors and proactively address lead hazards. Over the past several years, several school districts around the US have been collaborating with universities to develop AI based systems to help them identify at-risk students who are unlikely to finish high school on time. Finally, the University of Chicago has been working, as part of the White House Police Data Initiative, to identify officers who are at risk of adverse incidents early and accurately so supervisors can effectively target interventions.

Work in this area requires deep and sustained interaction and efforts between the target community and AI researchers, but there isn't a ready supply of trained AI researchers (or practitioners) who are familiar with the unique aspects of working on public welfare problems. Likewise, government and policymakers have little experience working directly with the research community. Finding funding mechanisms that bring both communities together to address local needs -- e.g. the NSF Data Hubs model -- is essential. Highlighting ongoing projects (and successes) to both raise awareness and to provide a roadmap is essential to growing this community. Platforms that are able to access, structure, and curate appropriate data sets do not exist.

Projects need to have a long-term structure, with appropriate intermediate goals, to avoid short-term fixes, or quick, but ephemeral, "feel-good" stories. Legal and regulatory hurdles including, access to data, and to populations to evaluate against, will require substantial investment of time, planning, and resources to effect. Creating a framework for ethical evaluation of costs and benefits must be established. Understanding the impact of innovations will require an understanding of the level of compliance, and possibly methods to manage or pivot solutions in response to perception, trust, and compliance of the target population.

There are several related technical challenges. Privacy issues, transparency and traceability of data collection and decision-making, and understanding of social context must be considered within the research context. Issues surrounding data bias and uncertainty have direct implications to fairness and the evaluation of the utility of possible decision paths. Related (government) organizational and (population) sociological constraints must also be considered. More technical problems include: 1) data analytics and machine learning models that are robust to systematic bias, missing data, and data heterogeneity; 2) the development of models or simulations that sufficiently predict to inform decision-making, and which also can then be adapted "closed-loop" as additional data is collected with time; 3) advanced models of decision-making and planning that incorporate social dynamics, resource constraints, and utility models for multiple actors; 4) consistent, cost-effective, and scalable models for measurement or data collection; and 5) methods for causal reasoning and explanation.

Some near term opportunities include: 1) tracking of location data and understanding how to better predict/deploy first responders, 2) using individual public transit and other transportation data (uber, bikeshare, etc.) to understand mobility patterns of people to understand gaps in transit (where they live - where they work - what services they need) and also to assess impact of policy changes; 3) better detection of women who may be at risk of adverse births to target human services programs and resources; 4) better detection of adults in danger of becoming homeless/incarcerated; 5) increase the number of kids who are performing at grade level by creating interventions that would influence and change behavior; and 6) enhancing access to services/food/health.

Longer-term opportunities will build on the establishment of a platform for evidence-based decision making by government informed by more detailed and nuanced models. For example, is it possible to predict the acceptance or engagement of the population to a particular policy change. Also, such models could move toward a "systems of systems" analysis where information about welfare impacts education impacts law enforcement impacts health. Achieving these ends will require methods to integrate multiple AI systems, and monitor, detect, diagnose, and adapt to multi-faceted population behaviors.

Cross-cutting Observations and Recommendations

To date, AI has typically focused around deploying narrow wedges of technology in narrow application areas. However, as we look across application spaces, we see a common thread of needs and approaches that are necessary to scale these "niche" approaches to address broad socio-technical themes. Common themes in this report and our discussions include: 1) improving data quality and availability; 2) supporting technology and policies that ensure individual privacy and data security; 3) mechanisms to promote collaboration (at development time) and adoption (at deployment time) of innovations; 4) mechanisms to ensure fairness, transparency, accountability, reliability of decision; 5) methods to accurately measure and assess the effect of a technology intervention over varying timescale; 6) long-term programs that train scientists in developing AI methods for complex socio-technical systems.

References

Belard A, Buchman T, Forsberg J, Potter BK, Dente CJ, Kirk A, Elster E. Precision diagnosis: a view of the clinical decision support systems(CDSS) landscape through the lens of critical care. J Clin Monit Comput. 2016 Feb 22.

Chow, Y., Tamar, A., Mannor, S., & Pavone, M. (2015). Risk-Sensitive and Robust Decision-Making: a CVaR Optimization Approach. In NIPS 2015.

Collins, F. Big Data Study Reveals Possible Subtypes of Type 2 Diabetes. NIH Director's Blog. 2015. https://directorsblog.nih.gov/2015/11/10/big-data-reveals-possible-subtypes-of-type-2- diabetes/

Commuting: "The Stress that Doesn't Pay". Psychology Today. (2015) https://www.psychologytoday.com/blog/urban-survival/201501/commuting-the-stressdoesnt-pay

Fang, F., Nguyen, T. H., Pickles, R., Lam, W. Y., Clements, G. R., An, B., Singh, A., Tambe, M., Lemieux, A. (2016). Deploying PAWS : Field Optimization of the Protection Assistant for Wildlife Security. In Proceedings of the Twenty-Eighth Innovative Applications of Artificial Intelligence Conference. AAAI.

Henry, K.E., Hager, D.N., Pronovost, P.J. and Saria, S., 2015. A targeted real-time early warning score (TREWScore) for septic shock. Science Translational Medicine, 7(299), pp.299ra122-299ra122.

Nicol, S., Fuller, R. A., Iwamura, T., & Chadès, I. (2015). Adapting environmental management to uncertain but inevitable change. Proceedings Royal Society B, 282(1808), 20142984. http://doi.org/10.1098/rspb.2014.2984

Sadilek, A., Kautz, H., DiPrete, L., Labus, B., Portman, E., Teitel, J., Silenzio, V. (2016).

Deploying nEmesis: Preventing Foodborne Illness by Data Mining Social Media. Association for the Advancement of Artificial Intelligence, 2016.
Saria, S., & Goldenberg, A. (2015). Subtyping: What it is and its role in precision medicine. Intelligent Systems, IEEE, 30(4), 70-75.

Schulam, P., & Saria, S. (2015). A Framework for Individualizing Predictions of Disease Trajectories by Exploiting Multi-Resolution Structure. In Advances in Neural Information Processing Systems (pp. 748-756).

Sheldon, D., Dilkina, B., Elmachtoub, A., Finseth, R., Sabharwal, A., Conrad, J., Finseth, R., Sabharwal, A., Conrad, J., Gomes, C., Shmoys, D., Allen, W., Amundsen, O. (2010). Maximizing the spread of cascades using network design. In Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (pp. 517–526).

Smith, S., Smart Infrastructure for Urban Mobility. Slides from the AI for Social Good Workshop. http://cra.org/ccc/wp-content/uploads/sites/2/2016/06/Stephen-Smith-AIslides.pdf

Van Hentenryck, P. Reinventing Mobility with Artificial Intelligence. Slides from the AI for Social Good Workshop. http://cra.org/ccc/wp-content/uploads/sites/2/2016/06/Pascalvan-Hentenryck-AI-slides.pdf