Carnegie Mellon University
Block Center for Technology and Society

A Policy Maker's Guide to Artificial Intelligence for State and Local Governments: **Reaching Safe, Effective and Equitable Scale**

Executive Summary

Artificial intelligence (AI) holds the potential to greatly improve local public service delivery for state and municipal governments. New AI systems can reduce costs and help ensure more people get the services they need. But in order for these systems to be both effective and fair, they need to be developed with community input and values in mind. The recent outpouring of concerns around algorithmic bias has caused some observers to believe AI-assisted public services delivery to be too great a risk.

The reality is AI systems are already in place and helping public servants provide services in cities across the country. Drawing from how Charlotte-Mecklenburg (North Carolina) improved community policing by using AI to identify 'at-risk' officers to how Allegheny County (Pennsylvania) is keeping kids safe through AI-assisted child protective services, as well as insights from Carnegie Mellon University research, this paper outlines strategies to deploy AI in local governments.

As more cities and states recognize the potential benefits for AI systems to improve local service delivery, the demand for these systems will grow. But without federal support, earlier adopters will primarily be wealthy, large, and tech-savvy cities. **In order to promote safe, effective and equitable adoption of proven AI systems across local and state governments, Congress and the Administration should establish a National AI Strategy for Local Public Services.** The Strategy would bring together the federal funding streams that support state and municipal governments to provide assistance in three areas: AI talent, digital infrastructure, and regulatory and standards development.

There are several factors that enable governments to take advantage of the benefits that AI can bring. First, massive amounts of data are now readily available. Second, advances in computing and storage technology enable efficient processing of this data. These systems enable training of the data at a scale not possible before. In addition, new algorithmic techniques are being developed by leading researchers that allow data to be processed and transferred faster than ever. And most important when applying AI to policy making: new approaches are emerging to integrate deep understanding of concepts such as bias, fairness, context, and general ethical principles in AI systems—from design to deployment. These approaches are essential to produce AI systems with societal benefits. This convergence leads to the development of new AI tools such as those described herein.

Martial Hebert
Dean, School of Computer Science

Introduction

Most articles and reports on artificial intelligence (AI) attempt to speculate the pace at which technological progress will be made and the associated opportunities and threats it presents. While thinking about the future is important, this report takes a different approach: it looks at the current state of AI and how its applications are already being successfully deployed in communities across the country. The current national debate regards the use of AI within the public sector with great skepticism—initially due to concerns over job automation, and most recently, with potential biases within the models. While these worries can certainly be warranted, it would be wrong to conclude these risks justify excluding the benefits of AI from cities and states.

AI is a platform technology that offers significant societal benefit. Some of the most exciting AI applications are already being enacted at the local level. To name just a few, they include:

- Reducing the risk of lead poisoning in children;
- Facilitating the early identification of disease epidemics[1] ;
- Curbing recidivism rates for people in need of mental health services;
- Improving educational outcomes for students at risk of not graduating from school on time;
- Bolstering police-community relations by identifying officers at risk of adverse incidents;
- Improving health and safety conditions in workplaces and in rental housing[2].

A significant risk, albeit less discussed, is the threat that these new opportunities remain the province of the wealthiest, largest, or most technologically sophisticated cities and states. Small and rural communities that do not have the necessary physical infrastructure and skilled workforce to take advantage of new technologies are at particular risk of being left behind. As with all new technology platforms, AI can deepen the digital divide between places and people; but, only if we let it. In order to avoid this outcome, the country needs a federally funded strategy to deploy the technical and physical infrastructure and support the talent needed to expand access to AI-assisted systems.

The goal of a National Strategy for AI-assisted public services should be the promotion of safe, effective, and equitable scale in order to not only even the playing field among communities, but also effectively address algorithmic bias by promoting best practices and shared models.

To argue by analogy, in the mid-2000s, the United States recognized that broadband was critical digital infrastructure and without substantial federal support, small and rural communities would be left behind. What followed was a National Broadband Strategy to help bring 21st century technology to all Americans. Of course, the democratization of high-speed internet has also brought cyberbullying, deepfakes, and foreign intervention of our elections into the living rooms of millions of Americans. But looking back, few would argue that we shouldn't have invested in broadband. If more cities and states are to enjoy the benefits that responsible AI systems can provide to their citizens, like broadband, then national investments in talent, digital infrastructure, and regulatory guidance need to be made now.

Ways Artificial Intelligence Can Support Public Services

The most obvious benefit to AI for public services is by simply helping programs meet their objectives better. New models can improve the efficiency and effectiveness of public services by both ensuring the right citizens receive benefits and the misallocation of resources is reduced. They can also help predict when and where future interventions will likely be needed and for whom, while identifying errors or mistakes that allow human operators to course correct. These models can also serve as a check on human bias or misinformation. This paper dives deeper into these topics in the case studies below, but beyond the ability to improve the quality of decision making, AI systems have several other benefits to policy making.

AI can support evidence-based policy making

State and local policy makers are often able to clearly articulate the goals and objectives of a given agency, but how do they know if the actions taken by those departments are actually achieving the prescribed outcomes? Fads, lack of resources, and other factors too often create an environment where policies and practices are created and deployed without a clear understanding of how to assess whether they actually work. In some instances, local governments will pursue the same policies for decades without actually knowing if they are effective. According to a review of over 15,000 national, state, and local programs, researchers at the London School of Economics found only 4.5 percent were adequately evaluated to ensure their effectiveness.[3]

Because artificial intelligence is by definition data-driven, algorithm-assisted approaches and practices are well-positioned to help state and municipal governments test, iterate, and refine programs. As with any policy, public officials may choose not to use the data to improve decision making. However, the fact that algorithms require robust data at the outset has already set these types of practices apart from many programs that are implemented without corresponding data. Machine learning models "learn" by iterating between historical data and new information, which is why it is important to integrate algorithms and human service providers together to ensure decisions reflect the programs' underlying intent.

AI requires communities and policy makers to explicitly define values

AI systems require public officials to define exactly what to optimize for and which mistakes are both financially and socially more detrimental. As such, ethical and societal values must be explicitly defined. While values are, of course, implied in the human decision-making process that exists today, they are not necessarily made explicit. When the implicit values that enter into human decision-making are biased and unfair, inequitable outcomes occur. This is a critical reason why thoughtful, well-structured AI systems can significantly augment or even overcome human bias.

For an AI system to function, these values need to be provided as a critical input. For example, a system that is recommending lending decisions will require the following decisions: 1) specify the differential costs of flagging someone as unlikely to pay back a loan and being wrong about it versus predicting that someone will pay back a loan and being wrong about it, and 2) specify those costs explicitly in the case of people who may be from different gender, race, income, or education-level groups. In the past, that may have happened implicitly, with high levels of variation across different human decision makers (loan officers in this case). With AI-assisted decision-making processes, we are forced to define these values explicitly at the genesis of the program.[4]

AI can improve community involvement in policy decisions

Even with values defined explicitly in AI systems, who makes those decisions is still determined by the policy making process. Who and how these values are defined cannot be left to the AI system developer or an arbitrary set of individuals who define those values in an AI algorithm, be it explicitly or implicitly. Even when the right value judgments are made, without community understanding and support, suspicion over AI can upend otherwise sound projects. Most failed AI-assisted public programs do so at least in part because the community and those affected or neglected by the program are not involved in the decision-making process.

A telling case is the water crisis in Flint, Michigan. In 2016, 15,000 homes in the city were exposed to dangerous levels of lead in their drinking water. Due to poor public records, the city couldn't identify which homes were at risk.[5] In 2017, a team of volunteer computer scientists designed a machine learning model to determine which homes were most likely to have lead pipes.[6] Using the model, workers inspected 18,883 homes and 6,288 had their pipes replaced—a 70 percent accuracy rate.[7] However, the city abandoned its use of the algorithm in 2018 due to community backlash, which resulted when homeowners who didn't receive inspections saw that the model's predictions instead chose their neighbors. Without sufficient explanation, the community lost trust in the algorithm. By 2018, the rate of inspections to lead pipe identification fell to less than 15 percent. Through greater community engagement, Flint has redeployed the algorithm, but with significant costs and lost time.

Citizens have mixed feelings about AI. They are rightfully suspicious when it comes to applying the technology to essential public services. Elected officials and civil servants can help to alleviate these concerns by involving the public in each step of the process. An important first step is to find simple ways to explain how an algorithm leads to decisions, the types of errors it may encounter, and the trade-offs the machine makes. The standard diagram used by computer scientists to do this is called a "confusion matrix." Though the name doesn't literally derive from human confusion, it is rather apt. On their own, these diagrams are extremely unclear to lay readers. Within a larger body of work around community perception of algorithmic bias, Carnegie Mellon University researchers have begun testing clearer ways of explaining a model.[8]

For example, suppose a municipality would like to introduce an AI system to assess a criminal defendant's likelihood of re-offending, and public officials would like to explain how the system treats two groups of defendants (Group A and B). They explain that the machine is usually accurate in its predictions. However, they admit that the model sometimes inaccurately predicts that a defendant will re-offend when they don't (false positive) and that some defendants won't re-offend when they do (false negative). Graph One below shows the standard confusion matrix, with technical terminology, while Graph Two illustrates a revised flow chart. Not surprisingly, the CMU team finds that by simply using common terms and expressing the direction of decisions, average citizens are significantly better able to understand the algorithm.

Graph One: Standard confusion matrix (with original terminologies)

		actual class	
		positive	negative
predicted	positive	true positive 1369	false positive 805
	negative	false negative 532	true negative 990

Group A

		actual class	
		positive	negative
predicted class	positive	true positive 505	false positive 349
	negative	false negative 461	true negative 1139

Group B

Source: Maria-Florina Balcan, Jason Hong, Ariel Procaccia and Hong Shen, Carnegie Mellon University, 2020.

Graph Two: Simplified flow chart with common terminology

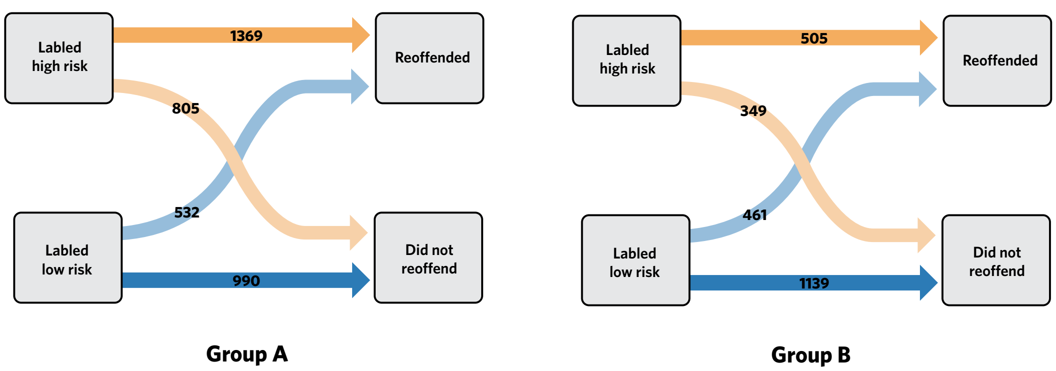

Source: Maria-Florina Balcan, Jason Hong, Ariel Procaccia and Hong Shen, Carnegie Mellon University, Forthcoming.

The point is, algorithms can be directed to make certain types of mistakes over others (more false positives or false negatives) and these decisions are not technical but ethical and social. Different communities will have opinions on how the AI system operates based on their values. Public officials need to be intentional about finding ways to explain the choices and trade off options of a given algorithm.

Existing Applications of AI for Public Services

Today, cities and states are addressing some of our country's biggest challenges by adopting new AI applications. While these models are not yet perfect—no machine or human models are—they have already proven to significantly support public servants and improve the quality of life for people within their jurisdiction.

Charlotte-Mecklenburg improves community policing through machine learning

The Challenge

The 2012 killing of 17-year-old Trayvon Martin, which led to the emergence of the Black Lives Matter national campaign, has heightened the national conversation around community policing and race. In hundreds of cities and towns across the country, the relationship between the police and the communities they serve is at best defined as mistrusting and hostile. Community trust of police is essential for neighborhood safety; for example, a lack of trust is associated with reduced reporting of violent crimes.

In any given year, eight to nine percent of officers will have an "adverse incident"—a catchall term to describe events that create hostility between the public and the police—ranging from being disrespectful to a victim's family, to racial profiling, to deadly shootings. For decades, police departments across the country have been interested in identifying officers that might be at risk of these behaviors. Few interventions have proven comprehensive.

Conventional Office Detection Method Prior to AI

In order to detect officers at risk of adverse events, many departments have deployed "Early Intervention Systems"(EIS). In 2007, over two-thirds of medium and large departments in the United States had an EIS in place.[9] While better than nothing, most EIS have not proven effective at consistently identifying the officers most likely to run into trouble. There are multiple reasons. First, the standard EIS is not data-driven, but instead relies on expert observation and intuition. Second, because most systems rely on a threshold, they cannot identify very high or low performing officers. Third, these systems lack customization—such as a midnight shift in a high-crime area compared to a weekday shift in a business district—and often rely on fairly obvious indicators, which make them easy to manipulate.

Charlotte-Mecklenburg's New Predictive Model

In 2015, Charlotte-Mecklenburg Police Department introduced a machine learning model that drew from dozens of variables, including complaints, uses of force, neighborhoods patrolled, etc., to provide a risk score for each officer. The specific score allows for more targeted training, counseling, and other interventions for officers who are at the highest risk of having an adverse incident. Since its inception, the model has increased the accuracy of identifying officers at risk by 12 percent (e.g. true positive) and reduced the number of officers targeted as at risk that do not exhibit worrysome behavior by 32 percent (e.g. false positives).[10]

This reduction in false positives has allowed the department to better allocate resources, reduce the burden on supervisors, and reduce unnecessary administrative work of officers who were not at risk. Officer buy-in has also increased. Machine learning models such as the one in CMPD also return control to the department, allowing its leaders to choose the right mix of accuracy and interpretability. Finally, machine learning approaches can be used to generate more comprehensive risk scores as opposed to a pure "yes/no" classification. In addition to being a better fit for the resource constraints faced by today's American police force, risk score systems can identify which officers are doing well as easily as they can pinpoint which are at risk. The department can use this information when assigning officers to partners or identifying best practices to incorporate into its training programs.[11]

According to Carnegie Mellon University's Rayid Ghani, who helped CMPD design the system, "a data-driven approach allows police departments to move from punitive interventions to an effective and preventative program designed to improve police-community relationships." Looking forward, the team hopes to develop dispatch-level models where an officer at a higher risk of an adverse incident for that dispatch can potentially be held back and a different officer, with a lower risk score, can be sent instead.

Dispatch-level machine learning is critically important. In 2015, a police officer in Texas pulled his weapon on children at a pool party. It was later determined that the officer had responded to two suicide calls earlier in his shift.[12] By keeping departments attuned to the mental states of their officers, better predictive models can help to ensure situations like these occur less often.

Pittsburgh keeps kids safe with predictive child services screening

The Challenge

In 2015, over 40 percent of the seven million children reported to protective services in the United States were screened out primarily based on the discretion of local screeners. During that same year, 1,670 children died as a result of abuse or neglect. While protecting children from neglect and abuse is a paramount priority to any local government, municipalities face the herculean task of not only responding to each call, but deciding which constitute legitimate threats.[13] The practice of deciding which calls merit further in-person investigation is left to each jurisdiction to follow local practices and policies, potentially leading to large variations in the way referrals are treated across the country.

At the same time, first-generation AI models have provided critical support to local screeners, improving efficiency and, most importantly, accuracy. However, these systems are often criticized for over-accounting for certain demographic groups. As these models are distrusted by many community members, their implementation remains a contentious issue.

Call Screening Prior to AI Support

While every municipality has its own specific policies, the child welfare system as a whole is universally responsible for responding to all cases where there is significant suspicion that the child is in present or impending danger. As a case in point, Allegheny County, Pennsylvania, has a population of 1.2 million people and receives roughly 18,000 allegations of maltreatment each year. The county's child protective hotline staff is responsible for screening and deciding which calls constitute an in-home visit. While staff has access to a significant amount of data, including historical information on public services (e.g. drug and alcohol services, homeless services) and individual and family data, it is challenging for county staff to efficiently access, review, and interpret all available records. Beyond the time required to scrutinize data, most municipalities have no means of ensuring that the most relevant available information is consistently used or weighted by staff when making hotline screening decisions.[14]

While much of the public debate around protective services revolves around the quality of service, the real problem often centers on whether or not the right families are being served. Consider: in Allegheny County, of the 18 cases where children were later killed or gravely injured between 2010 and 2014, eight, or 44 percent, had been screened out.[15]

Allegheny County's AI-Assisted Screening Tool

In August 2016, the Allegheny County Department of Human Services (DHS) implemented the Allegheny Family Screening Tool (AFST), a predictive risk modeling tool designed to improve child welfare call screening decisions.[16] The AFST was the result of a two-year process of exploration on how existing data could be used more effectively to improve decision making at the time of a child welfare referral. To generate the AFST scores, the system uses more than 100 predictive factors for each child on the referral.[17]

Describing one instructive incident by a New York Times article on Allegheny County's system, Dan Hurley notes: "Finding all that information about the mother, the three children, and their three fathers in the county's maze of databases would have taken [the screener] hours he did not have; call screeners are expected to render a decision on whether or not to open an investigation within an hour at most, and usually in half that time. Even then, he would have had no way of knowing which factors, or combinations of factors, are most predictive of future bad outcomes. The algorithm, however, searched the files and rendered its score in seconds."[18]

Carnegie Mellon University researchers are now building upon the success of the AFST by developing methods to ensure that such algorithms are both accurate and fair to all families. For example, one problem identified by Carnegie Mellon's Alexandra Chouldechova is predictions made from historically observed actions and interventions. In this historical view, families who are low-risk will have similar predicted risk to families who are high-risk but who receive supportive services (e.g., child care, employment support) and because of those services no longer exhibit high risk behavior. In other words, algorithms trained in this way will conflate families that are low risk without intervention with those who are low risk precisely because they are helped by system intervention. By better accounting for the effects of supportive services, the algorithm will have fewer "false negatives," or children that the machine thought were not at risk but are—and this translates into better identifying children who would likely benefit from supportive services.

Accounting for these factors could make the system fairer. As Chouldechova puts it, "we need to be careful in how we train and evaluate prediction models for use in decision-support settings. Many of the common pitfalls are statistically subtle but can have a significant impact on the tool's predictive utility and bias. Our work seeks to surface these pitfalls and to provide practitioners with the statistical methodologies to avoid them."

Currently, jurisdictions in California, Colorado, Oregon, and across Pennsylvania are all in various stages of developing and integrating artificial intelligence into their call screening processes. These improvements could significantly help how we protect children.

Supporting underserved communities while avoiding algorithmic bias

It is now well understood that algorithms can be susceptible to the same types of biases mirrored by human policy makers. This is particularly true when a model draws from historically-biased data, such as the relationship between neighborhoods (correlated with race) and crime.

Fear of algorithmic bias, coupled with the serious concern that many areas of public life have historically produced significant race and gender disparities, has led to backlash against the public use of AI among policy makers, advocates, and concerned citizens. Some advocates argue it is best to avoid the use of AI altogether in many public services because of the presence of such historical bias. This would be an unfortunate overaction that would deprive countless citizens of higher quality and less costly services.

It is important to recognize that AI can have a massive, positive social impact. However, we need to make sure that we put guidelines in place to maximize the chances of this positive impact while protecting people who have been traditionally marginalized in society and may be affected negatively by these new AI systems.[19] For example, algorithms that only optimize for efficiency can leave behind "hard cases" of people that actually need public support the most.

Policy makers should be concerned primarily with the policy outcomes of AI-assisted decision-making tools. That means more than just "fair algorithms." Even equitable algorithms can lead to unfair outcomes if the human decision of which intervention to pursue benefits some communities more than others. At the same time, it's possible for an algorithm that draws from historically-biased data to be fair with the right bias mitigation plans. As Carnegie Mellon's Rayid Ghani explains, "in some recent preliminary work we did with the Los Angeles City Attorney's office, we found that by careful consideration and analysis, we can mitigate the disparities that a potentially biased algorithm may create, and coupled with a tailored intervention strategy, the system has the potential to result in equitable criminal justice outcomes across racial groups."[20]

In the early days of "AI for Good," algorithms were too often treated as if the fact that they were derived from math implied that their decisions reflected unbiased, universal truths. Such overconfidence is unfounded. That said, it would be equally tragic if the pendulum swings too far in the other direction and we lack any confidence to leverage new technology to help our most vulnerable citizens.

Recommendations to scale AI-supported public services

The goal of this paper is to highlight how artificial intelligence is already helping state and local governments provide public services. Participating agencies are seeing improvements to the quality, cost and delivery speed of services. Just as AI is constantly changing, these public systems need to be maintained and improved. State and federal governments can support those maintenance and upgrade efforts. They also should help expand digital access to new AI systems to municipalities across the country.

In February 2019, the Trump Administration released its Executive Order on Maintaining American Leadership in AI.[21] It called for a number of important activities, including increased AI research and development and the number of AI trained workers. The order also called for the creation and improvement of national standards for algorithms. While all critical, the Executive Order did not comment on the role the federal government can play to ensure municipal and state governments have the resources needed to implement AI strategies.

In order to promote safe, effective and equitable adoption of proven AI systems across local and state governments, Congress and the Administration need to establish a National AI Strategy for Local Public Services. The Strategy would bring together the federal funding streams that support state and municipal governments, including relevant block grants to provide support in three areas: AI talent, digital infrastructure, and regulatory and standards support.

Increasing access to AI talent for state and local governments

Implementing AI systems requires people with deep knowledge of how these systems work. Currently, access to talent is the most significant barrier to expanding AI services to more cities and states. Particular federal strategies could:

- Establish an I-Corps like program to connect AI researchers at universities and national labs with municipalities to develop AI systems that reflect local realities.
- Create a Chief Artificial Intelligence Officer Academy to train the next generation of public sector CTOs and CIOs in new AI methods.
- Launch an AI for Good talent fund between local and national governments and philanthropy that would recruit, reward, and place computer science graduates into local governments for three years.

Investing in digital infrastructure

- Develop "system of systems" approach to local services that provides a framework and resources for interoperability of agency databases.
- Expand local and state cybersecurity capacity by identifying a typology of current readiness levels of different municipalities, ranging from small rural communities to major metropolitan areas. For example, municipalities could use support developing data maturity models, to self-assess current capacity.[22]
- Help local and state governments own and operate their own data systems, compared to giving their data to third party vendors. This also should include support for best practices around data governance.

Supporting regulatory guidance and standards

- Help cities develop "Citizens' Data Bill of Rights" to work with community members to structure guardrails to the use of AI based on citizen values.
- Expand the existing regulatory environment to account for AI-assisted decision-making. Instead of creating a new AI regulatory agency, the federal government should update the regulatory framework of existing relevant agencies including the SEC, FINRA, CFPB, FDA, FEC, FTC and FCC to ensure they apply to AI-assisted decision making.
- Procuring AI systems should include Key Requirements in the Request for Proposals (RFP) process to require proposers/bidders to include an explicit initial project period to identify what it would mean to have equitable outcomes as well as a continuous improvement plan to ensure the system is constantly monitored and improved.

Looking Forward

In many ways, Charlotte and Pittsburgh look like much of the country. While their local economies are growing they are not rich cities. And while both are home to several leading universities and technology companies, they are not tech meccas like Boston or San Francisco. And perhaps most important, like most cities, they must find new ways to address old issues that have only become more complex and contentious with time—without additional resources.

Where Charlotte and Pittsburgh begin to look different is how they have chosen to address these problems. They have thoughtfully leveraged new AI systems, not run from them. And each was successful because city leaders methodically and intentionally involved the community. Of course, neither were without setbacks, but ultimately these cities created systems that simply worked better than existing methods. And the stakes couldn't have been higher. In Pittsburgh, success meant fewer children abused and in Charlotte fewer deadly policy shooting.

Local governments share many of the same problems. They should also have the opportunity to share proven AI-assisted solutions.

About the author:

Scott Andes is the Executive Director of the Block Center for Technology and Society at Carnegie Mellon University. Please send any inquiries on this publication to sandes@andrew.cmu.edu.

About the cited Carnegie Mellon researchers:

Child welfare and creating fair algorithms:

Alexandra Chouldechova, *Assistant Professor of Statistics and Public Policy at Carnegie Mellon University's Heinz College of Information Systems and Public Policy*

Community policing in Charlotte, algorithmic bias and AI for Good:

Rayid Ghani, *Distinguished Career Professor in the Machine Learning Department and the Heinz College of Information Systems and Public Policy at Carnegie Mellon University*

Evaluating people's perception of fairness in machine learning:

Maria Florina Balcan, *Associate Professor of Computer Science*
Jason Hong, *Professor Human-Computer Interaction Institute*
Ariel Procaccia, *Gordon McKay Professor of Computer Science at Harvard University*
Hong Shen, *Systems Scientist, Human-Computer Interaction Institute*

Endnotes

1 Scott Andes, "Public trust in data could have helped China contain the coronavirus," *The Hill,* February 19,2020.

2 For more information on these examples, see: Rayid Ghani, "Equitable Algorithms: Examining Ways to Reduce AI Bias in Financial Services, Testimony before the Committee on Financial Services, U.S. House of Representatives, February 12, 2020.

3 Justin Parkhurst (2017), The politics of evidence: from evidence-based policy to good governance evidence. Routledge Studies in Governance and Public Policy. Routledge, UK.

4 Rayid Ghani, "Equitable Algorithms: Examining Ways to Reduce AI Bias in Financial Services, Testimony before the Committee on Financial Services, U.S. House of Representatives, February 12, 2020.

5 Melissa Denchak, "Flint Water Crisis: Everything You Need to Know," National Resources Defense Council, Washington, D.C. 2018.

6 Jacob Aberneth, Et. Al, "ActiveRemediation: The Search for Lead Pipes in Flint, Michigan," In KDD '18: The 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, August, 2018.

7 Alex Madrigal, "How a Feel-Good AI Story Went Wrong in Flint," *The Atlantic*, January 3, 2019.

8 Forthcoming work by Maria-Florina Balcan, Jason Hong, Ariel Procaccia, and Hong Shen, Carnegie Mellon University.

9 J. Manis, C. A. Archbold, and K. D. Hassell, "Exploring the impact of police officer education level on allegations of police misconduct," *International Journal of Police Science and Management*, vol. 10, no. 4, pp. 509–523, 2008.

10 Samuel Carton, et. Al. "Identifying Police Officers at Risk of Adverse Events", Forthcoming.

11 Ibid.

12 Ibid.

13 Dan Hurley "Can an Algorithm Tell When Kids Are in Danger?" *The New York Times*, January 2, 2018.

14 "Developing Predictive Risk Models to Support Child Maltreatment Hotline Screening Decisions," Allegheny County Department of Human Services, March 2017.

15 Dan Hurley "Can an Algorithm Tell When Kids Are in Danger?" 2018.

16 "Developing Predictive Risk Models to Support Child Maltreatment Hotline Screening Decisions," 2017.

17 "Frequently-Asked Questions on Allegheny Family Screening Tool (AFST), Allegheny County Department of Human Services, March 2017.

18 Dan Hurley "Can an Algorithm Tell When Kids Are in Danger?" *The New York Times*, January 2, 2018.

19 Rayid Ghani, "Equitable Algorithms: Examining Ways to Reduce AI Bias in Financial Services, Testimony before the Committee on Financial Services, U.S. House of Representatives, February 12, 2020.

20 Ibid.

21 "Executive Order on Maintaining American Leadership in Artificial Intelligence," White House Executive Order, February 2020.

22 For more information on Analytics Maturity Models, see Norman Reitter and Barry List, "INFORMS Analytics Maturity Model," INFORMS, Dec 2013.