Statement of Georgia Tourassi
Distinguished Scientist
Director, Health Data Sciences Institute
Oak Ridge National Laboratory

Before the
Committee on Science, Space, and Technology
U.S. House of Representatives
June 26, 2019

Hearing on Artificial Intelligence: Societal and Ethical Implications

Chairwoman Johnson, Ranking Member Lucas, and distinguished members of the Committee: Thank you for the opportunity to appear before you today. My name is Georgia Tourassi. I am a Distinguished Scientist in the Computing and Computational Sciences Directorate and the Director of the Health Data Sciences Institute of the U.S. Department of Energy's Oak Ridge National Laboratory (ORNL) in Oak Ridge, Tennessee. It is an honor to provide this testimony on the transformative role and ethics of Artificial Intelligence (AI) in healthcare and the role of the Department of Energy and its national laboratories in spearheading socially responsible AI science and responsible use of the federal data assets for AI innovation.

INTRODUCTION

As a formally trained physicist, engineer, and applied scientist with three decades of AI experience in biomedical discovery and healthcare delivery, both in academia and a national lab, I have observed the growth, subsequent stagnation, and latest renaissance of AI for health. These experiences have shaped my views on the AI topics you are considering today. In my role as director of the Health Data Sciences Institute and my experience as a computational scientist at ORNL, I have a broader perspective on the potential role of the Department of Energy's national laboratories in supporting an energy-efficient, socially responsible, and secure AI national strategy. I will begin with a brief overview of AI activities at Oak Ridge National Laboratory, exemplifying the impact of AI for our health. Specifically, I will highlight how the Oak Ridge National Laboratory supports the development of a thriving environment for AI advancement by

(i) providing responsible management of national data assets,
(ii) accelerating innovation, and
(iii) supporting a targeted, ethical, and socially diverse research and workforce development strategy.

With examples drawn from healthcare, I will discuss some of the societal implications of AI and the need for research and development (R&D) investments to ensure ethical use of the data and effective and energy-conscious AI technology. Lastly, I will discuss a vision of how the Department of Energy's national laboratories' AI resources, capabilities, and extensive experience can transform healthcare whilst promoting AI innovation, stakeholder engagement, and security of the nation's most sensitive data assets.

THE IMPACT OF AI FOR OUR HEALTH

The explosive growth of AI is driven by the convergence of big data, massive computational power, and novel algorithms. Together, these three pillars enable AI to accumulate, analyze, and automate the delivery of functional knowledge in many application domains.

The national laboratories are a remarkable asset for the nation. Over the past 75 years, they have consistently provided the science and technology needed to address national problems. As a DOE national laboratory and Federally Funded Research and Development Center, ORNL is equipped with exceptional resources in all three of the pillars of AI to support the Department's mission needs. Researchers at Oak Ridge and other national laboratories have a wealth of experience in building predictive models by using supercomputers to solve complex first-principles physics equations. With the steadily increasing power of today's supercomputers and the massive data sets that are becoming available in a variety of areas, we are now in a position to build the same types of predictive models by training these new models using AI. The leadership-class supercomputers at the DOE national laboratories are a unique asset for this training step.

Through the Department's Strategic Partnership Projects program, we have been given the opportunity to apply our resources to develop novel and promising AI solutions to several challenges in healthcare, as outlined below.

AI for modernization of national cancer surveillance

Supporting efforts to reduce the cancer burden in the United States, the national Surveillance, Epidemiology, and End Results (SEER) Program manages the collection of curated data from population-based cancer registries on cancer incidence, prevalence, survival, and associated health statistics for the advancement of public health. As part of the National Cancer Institute (NCI), the SEER program is tasked with supporting cancer research to improve the understanding of patient care and outcomes in the "real world" beyond the clinical trial setting.

NCI SEER partnered with DOE national laboratories to leverage their capabilities in high-performance computing and AI, with two goals in mind: to develop a more timely, comprehensive, scalable, and cost-effective cancer surveillance program, and to lay the foundation for an integrative data-driven approach to modeling cancer outcomes at population scale. By leveraging ORNL's exascale computing technologies, computational and AI algorithmic advances achieved through the DOE-NCI partnership provide the scientific framework to evaluate and inform personalized therapies, support the development of prospective diagnostics and treatments, and optimize population health outcomes.

The collaboration has had a direct impact on cancer surveillance over the first three years. Deployed AI algorithms have demonstrated high accuracy over 12 SEER cancer registries and millions of pathology reports received at the registries and processed with increased efficiency. The high accuracy and time efficiency of AI are expected to reduce the workload burden for cancer registrars while allowing them to realign their focus on abstracting additional complex variables (e.g., new cancer biomarkers, cancer recurrence) that are currently not possible using AI. Furthermore, by leveraging AI for “real time” cancer incidence reporting, we are moving towards real-time eligibility assessment of cancer patients for clinical trials. This is key to NCI’s goal of expanding the number and representativeness of patients in clinical trials and other research studies.

This effort to modernize the national cancer surveillance program exemplifies the benefits of a federally coordinated strategy to leverage AI, high-performance computing, and sensitive health data assets for real-world application. In addition, the specific application domain enables ORNL computational scientists to make fundamental contributions to the DOE mission with respect to preparing high-performance computing for the exascale as well as advancing the fundamental computational sciences of AI.

AI for intelligent clinical trials matching

ORNL recently participated in The Opportunity Project (TOP) Health Sprint. This was a 14-week effort sponsored by the US Census Bureau, coordinated by the US Department of Health and Human Services, and led by two Presidential Innovation Fellows. ORNL researchers showcased how AI can refine and advance the process of matching cancer patients to promising clinical trials.

Clinical trials have great potential in advancing the standard of care. However, matching patients with clinical trials remains a challenge, mostly due to the unstructured nature of eligibility criteria as well as the clinical documentation. The approach that we proposed leverages the power of AI technologies such as large-scale knowledge graphs and deep learning to bring together cancer registry data, medical ontologies, and clinical trials data to answer complex questions and provide real-time feedback for patients and clinicians on novel experimental treatments that are available to them.

AI for supporting breast cancer diagnosis

In an effort to reduce errors in the analyses of diagnostic images by health professionals, ORNL scientists have been working to understand and improve the cognitive processes involved in medical image interpretation. Our work has potential to improve health outcomes for the hundreds of thousands of American women affected by breast cancer each year. Our ORNL team found that analyses of radiographic breast images by expert radiologists were significantly influenced by context bias, or the radiologist's previous diagnostic experiences. So lab reserchers designed an experiment aimed at following the eye movements of radiologists at various skill levels to better understand the context bias involved in their individual interpretations of the images. Using the ORNL supercomputers, our team was able to rapidly train deep learning models necessary to make sense of large datasets of eye-tracking and medical imaging data. These findings will be critical in the future training of medical professionals to reduce errors in the interpretation of diagnostic imaging and will inform the future of human-AI interactions and personalized medical decision support and human augmentation going forward.

AI to improve the health of veterans: MVP-CHAMPION

ORNL is supporting the Veterans Health Administration in the analysis of its massive data resources, including the genomic data gathered through the VA's Million Veterans Program (MVP). The DOE-VA partnership, called the MVP-CHAMPION (Computational Health Analytics for Medical Precision to Improve Outcomes Now), began with a focus on the three health problems that affect the largest number of veterans: prostate cancer, cardiovascular disease, and suicide.

ORNL has brought its global leadership in computing and big data, as well as its demonstrated ability to analyze protected health information on a large scale, to the task of analyzing the VA's data stores. Our work has included the development of advanced algorithms and data mining techniques and of a novel data infrastructure that is consistent, regularly updated, flexible, and easily accessible. These efforts are helping VA and academic researchers tackle the country's most complex health challenges while advancing the state of the art in data analysis and computing, providing benefits to the VA, to DOE, and to the nation's veterans.

AI for Medicare and Medicaid: CMS

A few years ago, ORNL applied its resources to create a Knowledge Discovery Infrastructure (KDI) for the Centers for Medicare and Medicaid Services (CMS). This platform enabled comprehensive and longitudinal analysis of the extraordinarily large healthcare datasets maintained by CMS, which processes billions of claims each year, and identified patterns indicative of waste, fraud, and abuse. ORNL is now collaborating with the CMS Center for Program Integrity in an effort to develop advanced data analytics methods for detecting fraud.

THE HARD TRUTH OF AI

These examples demonstrate the potential of AI to revolutionize healthcare delivery. While machines won't replace physicians and nurses any time remotely soon, they do have enormous potential to assist health professionals and other decision makers with time-critical decisions for the smart delivery of healthcare. Furthermore, such studies will inform human-computer interactions going forward as we use AI to augment and improve human performance.

Still, the hope of AI must be tempered. AI is often exaggerated with hype and unrealistic expectations of universal benefits. This can lead those outside the AI research communities to believe that AI is the silver bullet that will solve any problem and overcome the known, and unknown, challenges facing us now and in the future. The reality of AI is that it is experiencing a renaissance in terms of capabilities, but still faces the necessary reality check that those capabilities have limitations and there are social implications if we fail to transparently address these limitations. With AI's great promise comes an even greater responsibility. As much positive impact as AI can provide to the nation, we must recognize the pitfalls and possible ill-intended uses of this powerful technology. In this section, I will address separately the ethical and societal implications related to (i) sensitive data management and use, and (ii) AI algorithmic development and practice in the real word, both in the context of healthcare.

The Ethics of Data: Access to large amounts of data is fundamental to AI. Furthermore, access to diverse and seemingly irrelevant data offers the most exciting opportunities in many application domains and certainly in health. For example, the adoption of electronic health records has created great opportunities for biomedical research to understand individual and population-level health outcomes over time. However, liberating and providing access to this data is both a technological and a policy challenge. The different facets of healthcare (clinical services data, pharmacy data, billing claims, insurance) are often siloed, in part due to policy and the intended use of the collected data. To create a richer picture, medical data cannot be siloed and must be combined with other data points, such as those providing context on a patient's living conditions, a well-known "gap" in the data collection system of the healthcare system that has substantial implications for care delivery. To address this gap, there is a big push in the field to include nontraditional datasets, such as socioeconomic data, socio-demographic statistics, environmental data, and even weather data, in predicting patient health trajectories. Broadly speaking, the biomedical research community is looking for ways to understand and describe the human holistically. At the same time, the sheer volume, variability, and sensitive nature of the personal data being collected require newer, extensive, secure, and sustainable computational infrastructure and algorithms.

Driving sustainable, secure, and energy-efficient infrastructure to handle increasing health data and AI computation: There is a pressing need to consider federal investments in centralizing national data assets in an environment that can provide the compute resources, a secure data management infrastructure, and scalable data analytics capabilities. The Department of Energy national laboratories are uniquely positioned to play that role for the nation, given the Department's long history of serving as the steward of large data infrastructures and of the nation's nuclear security enterprise. With their leading role in high-performance computing and their extensive data science and AI capabilities, the national laboratories could serve as a neutral entity, an honest broker for democratizing AI, while providing meaningful and responsible access to sensitive data assets and compute resources. Such an investment is critical to support the continuum of scientific discovery for effective domain-specific application.

Though AI is currently accelerating research across many fields and industries, an outstanding issue for sensitive domains, such as health and medicine, is how to preserve privacy while computing with shared sensitive medical data to obtain relevant insights. Such data includes personally identifiable information, personal health information, intellectual property, and other proprietary details, where potential leaks would have obvious adverse consequences. Removing personal identifiers and confidential details is insufficient, as an attacker can still make inferences to recover aspects of the missing data. Inference attacks can also jeopardize AI algorithms over shared data by targeting the shared AI model training process and the trained model itself. Indeed, serious threats are encountered in collective AI endeavors that aggregate data from different sources, since the most vulnerable source establishes the overall security level.

At ORNL, we fully understand the need to overcome these threats as we work to fully utilize the wealth of information in shared data, including data within the ORNL secure data environment, and to apply our high-performance computing resources for AI to make otherwise impossible scientific discoveries and technological breakthroughs in sensitive domains. This is an underdeveloped field of research in which R&D investments are well warranted to develop new solutions so that the community can responsibly and privately share sensitive data for aggregated analysis, including training shared AI models.

Aside from the data security concerns, the energy demands of digital infrastructures and compute- and data-hungry AI algorithms pose energy challenges. AI and other computing activities are projected to use over half of the world's energy by 2040.[1,2] We cannot easily predict how we can balance AI's energy demands with AI's demonstrated ability to guide management of our energy resources more efficiently. DOE's national laboratories are well positioned to lead hardware, software, and algorithmic innovations and deliver AI solutions that consume less energy, a challenge DOE is already working to address as part of the Exascale Computing Initiative. Energy-efficient AI is a key to securing our ability to provide sustainable and affordable solutions with benefits to the environment and our national and economic security.

Driving responsible use of sensitive data: The ethical collection and use of human data is an outstanding challenge for the research community. Although we all recognize the scientific value of human data, the debate over data ownership is ongoing

[1] Greenpeace, "Clicking Clean: Who is winning the race to build a green internet?" Washington, DC, 2017.

[2] Semiconductor Industry Association and Semiconductor Research Corporation, "Rebooting the IT revolution: A call to action," September 2015.

in terms of how best to balance the promise of transparent AI innovation with the risks of unethical data handling, intentional or unintentional privacy breaches, and adversarial data use by hostile or malicious actors.

As scientists, it is our responsibility and as a nation it is our obligation to invest in research and development and related policies to support socially responsible design and use of AI. Data is the new currency in all facets of life and most notably in healthcare, disrupting traditional technology transfer and business models as well as blurring the line between research use and commercial use of patient data. To maintain a strong ethical AI framework, we need to answer this fundamental question: Who owns the intellectual property of data-driven AI algorithms in healthcare? The patient? The medical center collecting the data by providing the healthcare services? Or the AI developer? Clearly, no single entity alone could deliver the breakthrough AI technology. We need a federally coordinated conversation involving all of the STEM sciences, social sciences, economics, law, public policy, and patient advocacy stakeholders. The conversation should also address how to revise and continuously update outdated legislation and regulations to address the emerging complexities of an otherwise exciting technological development.

The Ethics of AI Development and Deployment: On the technical front there are pressing questions on data quality, data bias, AI interpretability, AI vulnerability to adversarial use, and human-AI integration to augment and not inadvertently handicap the human. These topics should be driving our federal research and development priorities and investments. Still, one size does not fit all, and the best technical solutions, regulations, and policies will probably vary across application domains. Therefore, we should recognize that science depends on integration of basic and applied activities. We should invest equally in foundational and applied AI activities, in high-risk disruptive ideas and lower risk incremental endeavors, so that we can reap short-term gains while working towards more sustainable long-term benefits.

Driving responsible development of AI technology: AI developers will need to offer solutions that are not only “on-average” accurate but also offer a measure of trustworthiness at the individual decision level. The latter would require a detailed explanation of each individual decision AI makes, as well as deeper understanding of the conditions under which the technology is exceptionally successful or alarmingly flawed. To understand these intricate issues, scientists need to dive deeper into the possible sources of errors and uncertainty, including biases in the data collected to train the AI algorithms and implicit biases in the way AI is embedded with the human. Unless these issues are addressed and properly regulated by the appropriate accreditation bodies in healthcare (or any other relevant domain facing similar challenges), we will not see real-world value of AI outside the anecdotal studies.

Driving responsible deployment of AI technology: For example, one area of research that deserves attention is how to most effectively integrate AI technology with the human "in the loop." Although there is increasing fear that AI will eliminate jobs, for the most part AI tools have been shown to be effective for narrowly defined, repetitive tasks. AI will not take the jobs of humans in healthcare, at least not any time soon, but it will change those jobs by improving how humans work, making humans more accurate and more efficient. Ultimately, humans and machines will have to work well together. But this synergy won't happen organically, as past health AI experiences have demonstrated. It is important to train healthcare providers in how to use AI responsibly, how to remain vigilant avoiding mistakes of over-reliance when supported by AI, and how ultimately to be knowledgeable users of the technology.

The DOE national laboratories with the support of other federal agencies could provide a secure environment for objective benchmarking of sensitive datasets and AI algorithms against community consensus metrics to detect, monitor, and possibly correct dataset biases or inconsistent AI technology performance.

Driving the training of an inclusive and diverse AI workforce: With respect to addressing the data bias issue, aside from raising awareness and working towards algorithmic advances, an important first step is to acknowledge that socially responsible AI development and deployment starts with an inclusive and diverse AI workforce. When the community of AI developers both reflects the diversity present in the user community and is embedded in that user community from the start, we will be better positioned to anticipate unintended consequences in the real-world use of AI technology. Therefore, I advocate a multi-pronged approach throughout the AI lifecycle. First, we need to invest in foundational statistical and data sciences. During the development phase, scientists should promote a rigorous statistical framework to monitor for potential biases in the collected data. During the deployment phase, AI developers should implement rigorous quality control, monitoring AI performance across subgroups to confirm robust performance or identify performance gaps. We should work to communicate to AI users openly and clearly what they should expect across various settings, and we should educate users so that they are informed consumers of the technology. For example, responsible use of AI technology should become part of our mainstream digital education. We cannot anticipate every blind spot and we should not blame AI for learning from implicit biases in the data. Humans do too! It is unavoidable that we will generate AI code faster than we can deliver appropriate policies and regulations. By being inclusive, diverse, rigorous, and vigilant, we can mitigate many of the aforementioned risks. We will be able to democratically answer the ethical questions which will confront us as developers, marketers, users, and regulators.

At ORNL, we are particularly proud of our work with area community colleges and universities to develop the next generation of AI-savvy citizens. Every year we provide high school and college students with internship opportunities to learn AI and use the world's largest computer to address pressing biomedical challenges. In partnership with the University of Tennessee, we offer an interdisciplinary data science and engineering graduate program that is supporting the development of a new generation of highly diverse biomedical scientists with compute and data core competencies. This program could serve as a model for other institutions.

TRANSFORMING HEALTHCARE WITH AI: WHY DOE AND WHY NOW?

Our national healthcare system faces some pressing challenges. The average life expectancy in the United States, 79 years, has increased by more than 30 years over the past century.[3] Much of this progress is due to medical advances with treatment strategies against age-related global killers such as cardiovascular disease, cancer, and

[3] http://www.cdc.gov/nchs/fastats/life-expectancy.htm

stroke.[4] Still, medical research has been less successful at prolonging healthy life (i.e., health span). About 80% of Americans older than 65 live with at least one chronic condition, and 50% live with two.[3] As the aging population in the United States is growing rapidly, the incidence of age-related, costly, chronic conditions such as heart disease, cancer, diabetes, and Alzheimer's disease is reaching epidemic proportions. US healthcare spending already comprises far more of the national GDP than any other sector, including defense, education, energy, and transportation.[5] With annual total costs of age-related diseases expected to skyrocket, exceeding $1.5 trillion in the US by 2030 (considering heart disease,[6] Alzheimer's,[7] and cancer[8] alone), the nation has a pressing need to reduce the economic burden of population aging. Spending cuts and revenue increases cannot solve the increasing federal healthcare costs due to aging.[9] Prolonging our lifespan without prolonging our health span is financially unsustainable for our nation.

AI could offer powerful solutions to these challenges. By leveraging our federal data assets, computing capabilities, and AI, we can develop a strategic roadmap to extend health span and rein in healthcare costs by understanding the broad spectrum of all factors impacting well-being. For the first time, we are at the tipping point to map the human genome (i.e., genomic profile), phenome (i.e., physiologic status), and exposome (i.e., physical and social environment) in real time and across the human lifetime. This is an ambitious endeavor that requires transformative scientific advances in data storage, management, and analytics using massive amounts of heterogeneous health data.

Understanding the human genome-phenome-exposome interplay also demands a multidisciplinary team of biologists, physicists, chemists, engineers, mathematicians, computer scientists, and clinicians as well as the cooperation of key federal, academic, and industrial stakeholders. DOE and its national laboratories have the breadth of

[4] http://www.nejm.org/doi/full/10.1056/NEJMra1109345
[5] http://www.usgovernmentspending.com/year_spending_2017USbn_18bs2n#usgs302
[6] http://www.cdcfoundation.org/pr/2015/heart-disease-and-stroke-cost-america-nearly-1-billion-day-medical-costs-lost-productivity
[7] http://www.nejm.org/doi/pdf/10.1056/nejmsa1204629
[8] Yabroff, K. Robin, et al. "Economic burden of cancer in the United States: estimates, projections, and future research." Cancer Epidemiology Biomarkers & Prevention 20.10 (2011): 2006-2014.
[9] https://www.cbo.gov/publication/51580

science and engineering expertise, instrumentation tools and sensor technologies, big data science capabilities, and unique computing resources to tackle this grand challenge. Furthermore, DOE national laboratories are well suited to serve as “honest brokers” of health data to facilitate the complex interactions and cooperation among the various biomedical and healthcare delivery stakeholders.

The Human Genome Project, which was initiated by DOE, was a $3.6B federal investment through its completion in 2003.[10] The human genome-phenome-exposome mapping, a far more ambitious endeavor, will require a multiagency, multinational approach, and the required federal investment could be estimated at $20B in today’s dollars.

The national impact of this investment cannot be underestimated. Studies estimate that every federal dollar invested to map the human genome returned between $60 and $140 to the US economy. AI for health could have similar economic impact by generating public and private sector jobs, enabling the development of novel measurement and analytics technologies, and spearheading the health span revolution, a data-driven paradigm shift in human health sciences. Moreover, increasing the average health span by 2 to 3 years could reduce federal healthcare costs by as much as $7.1 trillion over 50 years while improving national well-being and extending workforce participation and productivity.[11]

We can draw another lesson from the Human Genome Project. From its earliest stages, it included a program to address the ethical, legal, and social implications (ELSI) of genomic research. This program, funded at 3 to 5% of the total project budget, became a model for ELSI programs around the world, and the National Human Genome Research Institute continues to fund ELSI research.[12] The program had a lasting impact on how the entire community, from basic researchers to drug companies to medical

[10] Hood, Leroy, and Lee Rowen. "The human genome project: big science transforms biology and medicine." Genome medicine 5.9 (2013): 1.
[11] Goldman, Dana P., et al. "Substantial health and economic returns from delayed aging may warrant a new focus for medical research." Health affairs 32.10 (2013): 1698-1705.
[12] https://plato.stanford.edu/entries/human-genome/

workers, used and handled genetic data. Continuing and expanding this research will help to ensure the responsible use of AI for health.

CONCLUDING REMARKS

The rapid adoption of AI has shown great promise across a wide array of domains such as healthcare, transportation, manufacturing, and cyber security, to name a few. It is becoming an integral part of our daily lives, offering solutions to problems as complex as drug discovery and as trivial as guessing what our next word will be when we text. AI can and will transform the national workforce. It can and will cause social and economic changes. Many of these changes will be positive, but AI can also reinforce and perpetuate harmful societal injustices. One thing is clear. As much as the science of AI is about difficult math and powerful computers, the impact of AI is also about the nation and its citizens. AI is a fundamental human endeavor.

Our nation faces a formidable set of challenges: ensuring our national security in a changing world; increasing the availability of clean, reliable, and affordable energy while protecting the environment; improving human health; and enhancing U.S. competitiveness in the global economy by fostering scientific leadership and encouraging innovation. AI is expected to offer solutions to many of these challenges, but the implications of this disruptive change cannot be underestimated. Close attention to the ethical, legal, and societal implications of AI will be required to ensure that its benefits are shared and its risks are managed and minimized.

A cohesive national plan for AI is an imperative need to secure the nation’s economic competitiveness and well-being. At the same time, we as a nation have the unique opportunity to create a well-defined, federally coordinated roadmap to engage in ethical AI development that delivers benefits across all private and public sectors, perhaps most notably in improving human health. The DOE national laboratories are uniquely equipped and positioned not only to make substantial contributions in addressing AI-driven opportunities and challenges, but also to support the execution of a national plan by enabling responsible use of the federal data assets for AI innovation. Although the views I offered in my statement are shaped by my scientific experiences in the healthcare space, they are shared lessons across other application domains.

Thank you again for the opportunity to testify. I welcome your questions on this important topic.