Federal Register Notice 86 FR 46278, https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Introduction

As a not-for-profit organization, The MITRE Corporation works in the public interest to tackle difficult problems that challenge the safety, stability, security, and well-being of our nation through the operation of multiple federally funded research and development centers and labs, and participation in public-private partnerships. Working across federal, state, and local governments—as well as industry and academia—gives MITRE a unique vantage point. MITRE works in the public interest to discover new possibilities, create unexpected opportunities, and lead by pioneering together for public good to bring innovative ideas into existence in areas such as artificial intelligence (AI), intuitive data science, quantum information science, health informatics, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resilience.

MITRE has a long history of partnering with federal agencies to apply the best elements of AI and machine learning (ML) while developing and supporting ethical guardrails to protect people and their personal data. Our team is committed to anticipating and solving future needs that are vital to the success and safety of the public and the country. In crafting this response, MITRE has drawn upon its rich experience in leading Government technology and policy solutions in healthcare, cybersecurity, AI assurance, social justice, identity management, systems engineering, and being an “honest broker” for data.

In the following pages, we offer thoughts on the National Artificial Intelligence Research Resource (NAIRR) roadmap elements (Question 1), including opportunities and goals (Element A), infrastructure (Element D), data (Element E), and security (Element F). We also provide input on NAIRR capabilities and services (Question 2), ethical considerations (Question 3), and democratization (Question 6). MITRE is eager to engage further with the NAIRR Task Force and we stand by ready to support this effort.

Response to Question 1

What options should the Task Force consider for any of roadmap elements A through I … and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

Element A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success

Recommendation – Extending beyond support to individual AI researchers, the NAIRR should aim to achieve: 1) community-based research collaboration, 2) hypothesis generation, and 3) a research marketplace.

MITRE has relevant experience stemming from developing infrastructure that supports the Million Veteran program for the Veterans Health Administration (VHA).[1] Like the NAIRR, the VHA infrastructure was designed to democratize access to data, tools, and computing resources to accelerate research. Although the two efforts aim to support different domains (AI vs. genomics), they aim to provide similar capabilities: provisioning and governing access to curated datasets, maintaining privacy, and making computing resources available. Based on the insights we gained while helping the VHA, we identify the following opportunities for the NAIRR.

Community-based Research Collaboration. Solving the wicked research problems of the next century will require community-based approaches. Many researchers will initially be interested in getting access to the NAIRR's resources so that they can train models germane to their individual research interests. However, the true potential of the NAIRR is more likely to be associated with designing mechanisms and incentives for catalyzing transdisciplinary research collaboration across NAIRR research activities. A community of researchers can be more effective in sharing data-curation methods and can model development approaches, positive as well as negative results, and even real-world performance information of AI models developed with NAIRR resources. There are several models of effective research collaboration that can be used to structure interaction mechanisms among researchers within the NAIRR—for example, see Trochim et al. (2008).[2]

Hypothesis Generation. As researchers use the NAIRR to tackle scientific and technological challenges that fall on similar lines of inquiry, there will be a significant opportunity to cumulate findings and establish a knowledge base that expands over time. Such cumulation may be automated to a certain degree if the computing environment can be instrumented to identify and extract actionable insights as models are trained. The resulting knowledge base can then be used to build intelligence to suggest what may be worthwhile to investigate next for a given line of inquiry. AI is being used for scientific discovery in many domains (e.g., see Daniels et al., 2021).[3] Developing such intelligence can be a fundamental research thrust that the NAIRR actively pursues.

Research Marketplace. Although the NAIRR should provide the necessary incentives for the open sharing and cumulation of data and resources, researchers may want to keep certain types of artifacts private—mainly due to intellectual property considerations. A "research marketplace" could be experimented with where researchers exchange algorithms, models, data, and perhaps even their own services through protocols they control. Emerging bartering approaches in collaborative networks (see Dalli et al., 2019)[4] may be of relevance to designing such a marketplace.

Element D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure

Recommendation – As a national resource supporting open AI research, it is vital for the NAIRR's infrastructure and shared resources to be secured on the network. The NAIRR should also promote research at the intersection of AI and cybersecurity – in particular, AI network defense.

MITRE has led several successful efforts over the years producing tools, datasets, and other resources aimed at enabling researchers and cybersecurity professionals to develop and evaluate network defense solutions that will be relevant to the NAIRR—see MITRE ATT&CK and related activities.[5] These efforts have demonstrated the value and efficiency of using common frameworks and concepts to foster collaboration and rapid progress for cybersecurity and network defense.

The NAIRR also affords the opportunity to advance research at the intersection of AI and cybersecurity. Our experience has taught us that designing, implementing, and maintaining network emulation tools and simulation packages requires upfront investment that can pose a barrier to researchers. The NAIRR could be resourced to provide the shared infrastructure and tools to a) emulate computer networks for the purpose of training AI network defenses at scale and b) simulate large collections of adversarial scenarios for domains such as autonomous vehicles—with the purpose of assessing mitigation of AI-targeted cyberattacks.

Element E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource

Recommendation – The NAIRR should incorporate the following pathways to trusted data sharing: 1) creating privacy-preserving derivatives, 2) making data resources interactive, and 3) setting up data enclaves within the NAIRR. The NAIRR will greatly increase its relevance and impact by incorporating these pathways, and to do so will require the enlistment of a trusted data broker with the experience, reputation, and unbiased independence.

One of the main objectives of the NAIRR is to provide access to large-scale, robust data sets conducive to a broad range of AI research. MITRE recommends the NAIRR incorporate pathways to trusted data sharing. In particular, the NAIRR should enable open research that leverages select sources of sensitive data in a manner that is responsible, reliable, and privacy preserving; thus, augmenting the NAIRR's data repository of public AI research data sets (which may be limited in number, size, and quality), with more operationally relevant data sets.

The Aviation Safety Information Analysis and Sharing (ASIAS) Program has worked through similar issues.[6] ASIAS is a partnership between the Federal Aviation Administration, MITRE, and stakeholders across the aviation industry. As a not-for-profit organization, MITRE serves as an "honest broker"—a trusted third-party role that fosters sharing of data amongst industry competitors, which would normally be difficult due to trust issues. The ASIAS model has resulted in a demonstrated ability to protect and appropriately handle sensitive data—a key factor in bringing new operators into ASIAS as the program expands and increases aviation safety for all stakeholders. Based on this experience and track record, MITRE recommends the NAIRR implement the following data pathways.

Creating Privacy-Preserving Derivatives. The NAIRR should selectively incorporate data from sensitive sources, containing sensitive information,[7] and through techniques such as de-identification and aggregation, create data set derivatives still useful for tackling open research problems while being privacy-preserving with sensitive data coded or redacted. MITRE has pioneered the use of cryptographic hashing to create non-reversible de-identification of sensitive fields in aviation data while retaining the ability to fuse multiple data sources for more context-aware analysis. For over a decade, MITRE has provided metrics, benchmarks, and analysis products to ASIAS stakeholders while ensuring that competitors' data remains de-identified, without tracing data back to sources.

Making Data Resources Interactive. The NAIRR should give researchers the ability to interact with (query, interrogate, and visualize) data within the NAIRR repository. Experience from the ASIAS Program shows that with careful construction of functions, data interactions can be successfully, responsibly performed directly on sensitive data. ASIAS provides interactive dashboards tailored to stakeholders to meet the needs of authorized users. Each capability has its own infrastructure for preparing the underlying data—some create aggregated extracts of the data to be shown, while others are based on data aggregation performed at the time of rendering. Each has its own security configurations to control data access.

Setting Up Data Enclaves. When the characteristics of a type of sensitive data are not conducive to de-identification or aggregation, then the NAIRR should incorporate the ability to create secured data enclaves that facilitate research by granting access to a controlled set of permitted users. MITRE has a proven track record of advancing the state of the art for solutions in automating the creation of secure analytical enclaves.[8] The NAIRR will benefit from such capabilities.

Element F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls

Recommendation – As a national resource supporting open AI research, the NAIRR should be resilient and protected against adversarial threats. The Adversarial Threat Landscape for AI Systems (MITRE ATLAS) knowledge base, tool set, and community can be of essential help.

The NAIRR must be cybersecure as a high-valued, heavily invested, networked, national resource. The security challenges are compounded by the way AI systems represent a unique and rapidly expanding attack surface with associated risks not addressed by traditional cybersecurity controls. These attacks can be a highly effective means of exfiltrating sensitive data, stealing intellectual property, or otherwise subverting the AI system (and in this case, the NAIRR) for malicious purposes, even when the AI models meet traditional assurance standards and are effectively secured from cyberattacks.

Consideration for adversarial attacks on the NAIRR should include issues such as data poisoning of hosted data sets, propagating unknown trojan horses and vulnerabilities within AI models and tools, and fostering research contributing to the development of technologies for nefarious applications. The source of threats to the NAIRR is also diverse, ranging from nation-state actors to cyberterrorists to overly adventurous students. NAIRR administrators must be proactive and not wait for a major crisis before they begin addressing such threats.

MITRE, utilizing input from Microsoft and a broad coalition of private sector companies, developed the Adversarial Threat Landscape for AI Systems (MITRE ATLAS).[9] ATLAS was developed by synthesizing real world case studies, voluntarily submitted by a wide range of industry partners, which detail real attacks conducted against AI systems. The ATLAS team used this to build a robust, common taxonomy of attack tactics and techniques that map to a broad range of contexts to empower security analysts across industry and within the government to detect, respond to, and remediate threats against AI systems.

Since its release in the Fall of 2020, ATLAS has rapidly impacted AI security across multiple industry verticals, with industry teams such as Microsoft, Bosch, Ant Financial Group, and Airbus testing their own AI systems using the ATLAS model. After using the model, these companies contributed the results of their tests as case studies to further improve the ATLAS knowledge base. Additionally, this collaboration resulted in Microsoft's release of a powerful open-source tool set called "CounterFit"[10] based on ATLAS, which gives companies and organizations who cannot afford dedicated AI security practitioners a robust ability to evaluate their own AI-enabled systems against known AI threats. The NAIRR should take full advantage of the growing ATLAS knowledge base and tool set, and the Task Force is encouraged to have administrators of the NAIRR participate and contribute to the ATLAS community.

Recommendation – The NAIRR should utilize federated identity technologies for its management of access controls.

MITRE has been working on the emerging trend of *decentralized technologies*[11] —in particular, *federated identity*[12]—which is the current direction for Identity & Access Management.[13] This is an ideal option for open platforms such as the NAIRR. Federated identity refers to linking an individual's electronic identity, authentication, and personal attributes across multiple web service endpoints using Single Sign-on. The website being accessed trusts the identity provider to validate their credentials using industry standard security protocols such as SAML 2.0, OpenID Connect, and OAuth 2.0. The current work MITRE is engaging in is related to the ID.me platform, which uses a federated identity model for the Login.gov identity and authentication service offering.

Response to Question 2

Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Recommendation – The NAIRR should provide the following data and privacy-preserving capabilities—synthetic data generation and federated learning.

The NAIRR can benefit from lessons learned from MITRE's work in facilitating research on health data. Health data is very tricky. Legally, it cannot be used for non-treatment purposes without patient consent with limited exceptions under carefully controlled circumstances.[14] If patient data is shared, subsequently exposed, the potential penalties are severe, and include criminal penalties. Health systems that possess patient data are very hesitant to share it, making it unlikely that a government agency or contractor will be able to amass health data for study purposes. However, there are ways to make sensitive data like health data more accessible, and the NAIRR should provide these as services and capabilities.[15]

Synthetic Data Generation. In addition to hosting data sets, the NAIRR should provide tools so researchers can augment their experimentation by generating their own synthetic data at scale. MITRE has extensive expertise in synthetic data generation. Its Synthea tool "is an open-source, synthetic patient generator that models the medical history of synthetic patients," providing "high-quality, synthetic, realistic, but not real, patient data and associated health records covering every aspect of healthcare. The resulting data is free from cost, privacy, and security restrictions, enabling research with Health IT data that is otherwise legally or practically unavailable."[16] With the NAIRR providing synthetic data generating services such as Synthea, researchers will have access to a scalable source of freely, unencumbered data; however, there are trade-offs to using synthetic data over real data. Synthetic data tends to be cleaner, more homogenous, and more structured than real data. Neural networks trained on synthetic data in the lab are at risk of not performing as well on real-world live data. Nonetheless, NAIRR researchers would benefit from on-demand supplies of scalable and tailorable synthetic data.[17]

Federated Learning. The NAIRR should support federated learning as an alternative to only training AI models monolithically on a centralized repository of data—particularly when facilitating research on more sensitive types of data like health data.[18] Federated learning distributes the AI model training such that the training data stays behind firewalls, providing greater security controls for data owners and stewards. It does require the different data holdings to use a standardized data model. This is an approach used by Observational Health Data Sciences and Informatics Community and several similar research networks.[19]

Recommendation – The NAIRR should host community challenge problems of national importance. For example, orchestrating solutions using AI to protect critical infrastructure based on the NAIRR providing a robust modeling and simulation (M&S) engine.

Protecting critical infrastructure—particularly public water infrastructure—is a national priority for the White House and the Hill. It is called out in the $1 trillion infrastructure bill and the Drinking Water and Wastewater Infrastructure Act. MITRE is actively working to support Government sponsors in the application of AI to protect critical infrastructure with an initial focus on water treatment plants.[20] This involves creating a computing infrastructure that supports collaborative experience and access to real-world modeling and simulation (M&S) objects to augment human intelligence and decision making through applied AI.

MITRE recommends the NAIRR provide M&S of critical infrastructure systems that facilitate the broad testing, evaluation, and development of AI-driven analysis, insight, and learning. Critical infrastructure M&S would encompass key aspects of urban environments where water, food, energy, health, finance, security, transportation, communication, and natural systems converge into complex interrelationships and communities. Many of these M&S objects communicate through a growing internet of things network, where AI-deployed analytics are becoming increasingly important for automation and augmentation of human decision making.

The goal is to generate realistic environments where monitoring, management, security, operational efficiencies, planning, and various other AI augmentation could be creatively and efficiently developed. For AI augmentation to be successful, it is vital for the NAIRR to facilitate and provide sufficient amounts of detailed, labeled data for training AI critical infrastructure approaches. Providing sufficient computing and M&S access will enable underserved AI critical infrastructure collaboration, analysis, and deployment. Further, MITRE recommends the NAIRR prioritize a standards-based, open-federation that emphasizes the joint use of M&S and AI pattern-detection approaches to enable the utilization of public and private digital resources.

Response to Question 3

How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Recommendation – The Task Force should consider establishing an AI review board (AIRB) process for the NAIRR, similar to Institutional Review Boards (IRBs), to promote equity and fairness. The NAIRR should promote research of mechanisms for accountability, transparency, and operational monitoring of AI technologies.

The challenges when developing AI are diverse and can be both technical and social in nature. As a result, no one person or discipline can singlehandedly “un-bias” AI or make AI ethical. The problem is especially relevant in AI where design teams tend to share many attributes (i.e., similar education and degrees, life experiences, and cultural backgrounds). If AI developers do not actively work to incorporate other valid perspectives into the development process, we risk having the AI reflect limited perspectives about how the technology will be used and by whom.

One potential avenue to promote diverse stakeholder involvement in AI research and development (R&D) within the NAIRR is to establish an AI review board (AIRB) that functions similarly to Institutional Review Boards (IRB) for human subjects’ research.[21] Such review/assessment should also consider baseline criteria for acceptable performance and risk. Baseline performance criteria should be both mathematical and contextual, and criteria should include the perspectives of all affected stakeholders. Risk assessment criteria should guide decisions about the AI’s suitability to a given application domain or intended use, including the level of clarity that different stakeholders require; risk criteria specific to those groups with greater needs; and guidance for higher-stakes cases when legality, ethics, or potential impact are of concern. These AIRBs should be composed of ethicists, subject matter experts, and representatives from communities who will be affected by the deployed model. In addition to "go/no go" research decisions, the AIRB should specify checks and balances with safeguards throughout the lifecycle of the AI technology to identify and mitigate biases/risks. Reviews will need to include data sheets,[22] AI model cards,[23] and impact assessments.

The NAIRR should also be used to conduct research into managing biases and risks with AI technologies by studying mechanisms and establishing best practices for accountability, transparency, and responsibly monitoring the impact of AI technologies once deployed in operational use.

Recommendation – The NAIRR Task Force should develop partnerships with industrial organizations that develop massive AI models affording AI researchers access necessary to study issues of equity, fairness, bias, accountability, etc.

Some of the most significant achievements in AI over the last few years have required massively large investments beyond the reach of most research institutions. For example, GPT-3, a language model developed by OpenAI, has 175 billion parameters, and required several thousand petaflop/s-days (10^{15} neural net operations per second for one day) of computation during training.[24] Scholars have expressed multiple concerns about the size and opacity of such models.[25] One concern, relevant to the NAIRR's goals, is that access to infrastructure to develop, apply, and evaluate such models is currently limited. For example, access to systems enabled by GPT-3 is currently moderated by OpenAI and GitHub who can withdraw access without public scrutiny or recourse.[26] This is an issue because researchers are not currently able to sufficiently research the limitations of these one-of-a-kind models nor sufficiently study the implications of their use. Specifically, researchers should have the ability to develop an understanding about how to prevent the misuse of such technologies, for example, by resourceful actors. Partnerships are recommended between the NAIRR and industrial organizations (that develop these massively opaque AI models) such that AI researchers may have the access necessary to study issues of equity, fairness, bias, accountability, etc. The Task Force should consider developing initiatives that foster access for researchers under the NAIRR umbrella to these large industry models and related technologies.

Response to Question 6

Where do you see limitations in the ability of the NAIRR (National AI Research Resource) to democratize access to AI R&D? And how could these limitations be overcome?

Recommendation – the Task Force can accelerate democratization of AI R&D by leveraging the NAIRR to: 1) overcome a lack of educator training in AI and shared educational resources; 2) provide AI training content that is interactive, easy to find, modify, and share; 3) supply access to scalable cloud-based computing, equitable training data sets, and academic journals, all easily accessible over the Internet; and 4) promote an educational culture with greater risk tolerance and flexibility, creating structures and incentives for more educators to introduce AI into their curricula and classrooms.

Democratizing access to AI research and development must include equitable access and awareness, training and support, and adequate resources within the NAIRR. For people to gain benefit from access to this infrastructure, they must first be provided equitable access and exposure to knowledge about AI itself and to realize its benefits and application. MITRE believes this exposure should occur during the formative education years and independent of one's field of study.[27] We believe this vision is limited in at least four areas:[28]

Lack of training. Lack of training on AI concepts and applications can lead to hesitation and lower confidence among educators to adopt AI material in the classroom. To overcome the tentativeness of teaching unknown material, educators across disciplines must be engaged directly and be provided with foundational training targeted towards specific competencies. Additionally, educators acting in isolation are limited by a lack of support and content sharing. The NAIRR's shared research infrastructure should encourage community support and content-sharing across all levels of the U.S. education system. Specifically, the NAIRR can help alleviate barriers to communication and support by providing a platform that encourages inter-institutional data and code sharing and promotes streamlined communication between learners, educators, and researchers.

Quality educational content. Currently, there is a lack of freely available AI/ML content that is interactive, easy to modify, and appropriately challenging. This means that content should be accompanied by supplementary material suggestions and lesson guidance for educators, allowing them to scaffold lessons by altering lesson vocabulary, contextual reading and discussion questions, and assignment difficulty as needed. Educators are also time constrained in lesson research, lesson planning, and mapping new lesson content to their curricula. Thus, the NAIRR's shared research infrastructure should serve as an environment in which educators and researchers can quickly research, obtain, modify, create, and share existing lessons that meet these criteria. The NAIRR can provide the capability to freely share modified material and additional lesson guidance. This will reduce the time and expertise needed for quality educational content, tailored to the subject area of interest, while encouraging community-based learning models.

Technical resources available. Students, educators, and researchers are limited by access to technical resources including analytic environments, high quality data, and required technology and hardware. For instance, socioeconomically disadvantaged students may be limited to engaging with content available only on a mobile phone through a free cloud-based analytics platform, while other institutions may conduct research using costly analytics software and leverage access to academic journal publications requiring paid subscriptions or other resources. These limitations could be alleviated by providing an analytics infrastructure that is web-based, accessible on mobile phones, and integrates with, or is centered around free or open-source tools and training data. Additionally, the equitability of training data should be considered in lesson development and research, especially in natural language processing applications, as training data sets are often derived from inequitable sources and can result in the perpetuation of biases.

Risk tolerance and flexibility. Educators have uneven levels of systematic incentives and flexibility that allow for them to put in the necessary time and energy to learn, use, and to build AI educational lessons for students. There is an abundance of interest in widening opportunities for students in data science and AI capabilities, but without the necessary top cover, flexibility in curriculum development, easing of existing time demands, and the establishment of personal growth opportunities and recognition, educators cannot easily participate in programs that might otherwise allow for them to introduce AI into their learning environments. Creating these structures, incentives, and flexibility for educators will further democratize access to AI by enabling more educators to participate in its development and offer educational opportunities to more students.

[1] Million Veteran program. https://www.research.va.gov/mvp/. Accessed September 27, 2021.
[2] Trochim, W. M., Marcus, S. E., Mâsse, L. C., Moser, R. P., & Weld, P. C. (2008). The evaluation of large research initiatives: A participatory integrative mixed-methods approach. American Journal of Evaluation, 29(1), 8–28. doi:10.1177/1098214007309280
[3] Daniels, M., Toney, A., Flagg, M., and Yang, C. (2021). Machine Intelligence for Scientific Discovery and Engineering Invention. Center for Security and Emerging Technology. https://doi.org/10.51593/20200099
[4] Dalli, D., & Fortezza, F. (2019). The new face of bartering in collaborative networks: The case of Italy's most popular bartering website. In Handbook of the Sharing Economy. Edward Elgar Publishing.
[5] MITRE ATT&CK is a globally accessible knowledge base of adversary tactics and techniques based on real-world observations; see https://attack.mitre.org/.
[6] ASIAS Stakeholders. https://portal.asias.aero/. Accessed September 27, 2021.
[7] Examples of sensitive information are personally identifiable information (PII) and protected health information (PHI).
[8] See MITRE Patent No. 10,795,709 – Systems and Method for Deploying, Securing, and Maintaining Computer-Based Analytic Environments.
[9] Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS). https://atlas.mitre.org/. Accessed September 27, 2021.
[10] Pearce, W., Shankar, R. and Kumar, S. AI security risk assessment using Counterfit. May 2021. https://www.microsoft.com/security/blog/2021/05/03/ai-security-risk-assessment-using-counterfit/. Accessed September 27, 2021.
[11] Anderson, M. Exploring Decentralization: Blockchain Technology and Complex Coordination. https://jods.mitpress.mit.edu/pub/7vxemtm3/release/2. Accessed September 27, 2021.
[12] How Federated Identity Can Increase access to Services and Benefits Online. ID.me. https://insights.id.me/wp-content/uploads/2020/06/How-Federated-Identity-Can-Increase-Access-to-Services-and-Benefits-Online-1.pdf
[13] IdAM in a Nutshell. https://public.cyber.mil/idam/idam-in-a-nutshell/. https://public.cyber.mil/idam/idam-in-a-nutshell/.
[14] Limited Data Set (LDS) Files. https://www.cms.gov/Research-Statistics-Data-and-Systems/Files-for-Order/LimitedDataSets Accessed September 27, 2021.
[15] Another benefit of synthetic data generation is to extrapolate data values that are not routinely collected or that could adversely impact individuals if asked, such as religious affiliation or LGBTQ+ status. The NAIRR can also benefit from MITRE's work with the "All of Us" research program on ensuring representative data. (See https://allofus.nih.gov/.)
[16] Synthea Empowers Data-Driven Health IT. https://synthetichealth.github.io/synthea/#about-landing. Accessed September 27, 2021.
[17] There are other options besides Synthea which you can find by searching on "synthetic health data."
[18] M.J. Sheller et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci Rep* 10, 12598 (2020). https://doi.org/10.1038/s41598-020-69250-1
[19] OHDSI. https://ohdsi.org/. Accessed September 27, 2021.
[20] MITRE's own work in this area is called MITRECity.
[21] For Office for Human Research Protections' *IRB Guidebook* see http://wayback.archive-it.org/org-745/20150930181805/http://www.hhs.gov/ohrp/archive/irb/irb_guidebook.htm.
[22] For example, see https://www.microsoft.com/en-us/research/project/datasheets-for-datasets/.
[23] For example, see https://modelcards.withgoogle.com/about.
[24] T. B. Brown et al., Language Models are Few-Shot Learners. 2020.
[25] See for example, E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" in Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 2021, pp. 610–623.
[26] Quach, K. A Developer built an AI chatbot using GPT-3 that helped a man speak again to his late fiancée. OpenAI shut it down. September 2021. https://www.theregister.com/2021/09/08/project_december_openai_gpt_3/. Accessed September 27, 2021.
[27] To this end, MITRE established Generation AI (Gen AI)—a consortium of faculty and students—preparing educators today (and tomorrow's workforce) to be creative, ethical problem solvers competent in the application of AI technologies. Gen AI is empowering teachers across the country through sharing lesson plans, curated data, and other classroom materials.
[28] These recommendations have been informed by MITRE's work on its Social Justice Platform, which "provides resources, data, tools, and frameworks that empower decision makers to create and sustain equitable solutions that bring positive change for a more just society." See https://sjp.mitre.org/.