

Response to the Request for Information on Enhancing AI Integration in Human and Health Services

Notice ID: ACF-AI-RFI-2024
June 30, 2024

Submitted to:

Office of the Chief Technology Officer
Administration for Children and Families
Department of Health and Human Services
330 C Street, SW
Washington, DC 20201
Attention: Kevin Duvall
Chief Technology Officer/
Acting Chief Artificial Intelligence Officer

Submitted by:

Mathematica
1100 First Street, NE, 12th Floor
Washington, DC 20002-4221
Phone: (202) 484-9220
Fax: (609) 228-4958

This page has been left blank for double-sided copying.

Contents

Introduction	1
Item-by-item responses.....	2
1. What AI offerings do you provide to customers in the human services space?	2
2. What is the level of AI literacy you generally perceive among your human services customers (prospective and existing)?	3
3. How is AI already augmenting human efforts in the outlined transformational areas (or any other human service area you believe is ripe for transformation with AI)?	4
4. Share recommendations you have for leveraging AI to advance human services mission outcomes.....	5
5. How should ACF and/or HHS incentivize innovative use of AI with appropriate safeguards?	6

Introduction

Mathematica is a nonpartisan research and data analytics organization with a mission to improve public well-being. Our staff have worked in close partnership with the Administration for Children and Families (ACF) for nearly 30 years. We currently support more than 20 active contracts for ACF that give us first-hand knowledge of ACF's mission, agencies, and programs, as well as the families these programs serve. Our deep, long-standing partnerships with ACF position us well to see both the enormous opportunities and the risks to public well-being presented by the development and deployment of artificial intelligence (AI). We welcome the chance to support ACF in considering the ways policy, research, training, and tailored technical assistance might support those opportunities and mitigate those risks.

We especially appreciate ACF's interest in having this conversation now: the pace of AI development and deployment is rapid, and without rigorous equity, governance, privacy, and security controls, as well as research and evidence on the impacts of AI, these advances could result in more harm than good. The following response draws on the broad and deep expertise of Mathematica staff across disciplines and policy domains to address many of ACF's critical questions about using AI to enhance operations across human and health services.

Item-by-item responses

1. What AI offerings do you provide to customers in the human services space?

AI is a transformational set of technologies encompassing software and statistical models that can learn, make recommendations or decisions, and communicate in ways that mimic human intelligence. Recent attention has focused on generative AI, which enables the creation of new text, media, and other outputs based on learned patterns. We view AI from an encompassing lens to capture both generative AI and traditional AI technologies, such as predictive analytics and natural language processing. This viewpoint is grounded in a recognition that individual AI technologies have distinct advantages and drawbacks and are best suited to different use cases across the human services landscape.

In our long history of working with ACF, we have seen how technical innovations can help address challenges, deliver results, and improve well-being through evidence. We are taking the same forward-thinking perspective in our approach to AI. Examples of our recent data-driven work with AI include the following:

- Using machine learning techniques to [identify promising workshop patterns](#) in healthy marriage and responsible fatherhood programs.
- Optimizing neural networks and other predictive algorithms to [identify students at risk academically](#) and inform development of early-warning academic systems.
- Employing machine learning and natural language processing techniques to [streamline tracking of occupational achievements](#) for the nation's doctoral scientists and engineers.
- Building agent-based models to [predict the spread of COVID-19 in schools](#) under various operational conditions.
- Developing a novel wastewater predictive modeling approach to [reliably forecast surges in communicable diseases](#).

In these projects and all other work involving AI, our approach to AI incorporates the following principles:

- **Responsible data stewardship.** Mathematica grounds its approach to AI in a foundational commitment to trusted data stewardship. Effective use of cutting-edge AI technologies must happen in combination with deep subject matter expertise, attention to the policy landscape, and rigorous data governance standards that safeguard data privacy and security while accounting for potential bias. As a mission-driven organization trusted for our ability to turn high-quality data into reliable insights, we know AI is only as good as the data behind it. We build security and reliability protocols in our use of AI—so that as the technology evolves, our approach does, too.
- **Promoting equity and well-being.** Our projects have shown AI can be a powerful tool in supporting decisions and augmenting services that promote public well-being in general and equity in particular. Our experience with AI tools has also informed us about many of the associated risks, including the ways poorly designed or implemented AI systems might exacerbate inequities. At the start of every new project, Mathematica screens for potential high-risk or high-impact applications of AI that require further oversight or careful consideration of methods to reduce potential bias. For example, Mathematica is incorporating a [fair machine](#)

[learning framework](#) to develop models that predict long COVID, to mitigate algorithmic bias and rigorously assess whether the model's performance is distributed equitably across racial and ethnic groups. Just like any other decision-making tool, how AI is deployed can promote either equitable or inequitable outcomes.

- **Transparency in model development, evaluation, and interpretation.** Transparency is critical to building trust in AI-backed solutions, especially to promote equity. A human-centered design approach that incorporates the voices of the data subjects as well as those who will use the AI solution as a decision support tool is essential. Mathematica followed this approach to predict unplanned hospital admissions and nursing home admissions for the [Centers for Medicare & Medicaid Services AI Health Outcomes Challenge](#). Working with our clinical and patient advocate partners, we used human-centered design to develop, test, and calibrate the model and address concerns over its utility and accessibility.
- **Quantifying uncertainty in model estimates and providing practical and interpretable findings.** [Bayesian statistical methods](#) enable nuanced discussions of model estimates and help contextualize findings with probabilistic statements. To further these goals, Mathematica developed [BASIE](#), an evidence-based Bayesian framework to interpret impact estimates. Building on these methods, Mathematica recently used Bayesian modeling to [reinterpret subgroup impact estimates](#) of experimental evaluations of employment programs for recipients of Temporary Assistance for Needy Families (TANF) and related populations, providing guidance for evaluators to deploy Bayesian methods for single- and multi-evaluation designs. We have also used Bayesian modeling [to stabilize measures of school performance](#), improving the reliability of performance measures and enabling state and local education agencies to provide support to the schools and students that need it most.

2. What is the level of AI literacy you generally perceive among your human services customers (prospective and existing)?

At Mathematica, the human services clients we work and partner with broadly include government agencies, practitioners and service providers, and the families these programs serve. We expect these groups will vary widely in their attitudes toward AI and in their AI literacy. We view AI literacy as a continuum, whereby people build foundational knowledge that enables them to interpret, evaluate, and use AI technologies effectively to support their decision-making processes. We believe many human services providers are currently in the earlier stages of this continuum, still accumulating baseline knowledge on AI use cases and safeguards.

Understanding the variance in underlying knowledge and also in opinions on the usage of AI—and what factors most influence AI knowledge and opinions—will be crucial to developing and using AI technology responsibly and responsively. AI should be viewed as an additive technology that supports human decision making by helping place relevant and practical insights at the fingertips of agency staff, practitioners, and policymakers. It enables them to use domain expertise and a deep understanding of contextual factors to make better decisions in less time. For example, child welfare case workers might use an AI-based recommendation engine to curate a list of possible local service providers for a client referral. The value of the AI tool is to streamline the search and data collection process to generate a list of service providers; the case manager can then dedicate more time to the referral decision, itself—applying their broad knowledge of child welfare systems and unique understanding of the client's needs.

Use cases that connect AI tools with real-world value for frontline human services program staff are powerful resources in helping build understanding of the capabilities of these technologies and promoting AI literacy. To that end, we have two recommendations for human services agencies seeking to expand AI literacy:

- **Invest in translational media** to inform agency decision makers—as well as service providers and program participants—about AI’s potential benefits, risks, and safety. Increased collaboration between human services agencies, researchers, technical assistance providers, policymakers, and the technology sector could promote translational work. Translational media (such as [infographics](#), [blog posts](#), [podcasts](#), [webinars](#), and [conferences](#)) can play a key role in clarifying what AI systems in human services look like in practice, providing agency leaders and policymakers with the knowledge to collaborate competently with researchers and technologists to develop reasonable goals and feasible standards related to AI.
- **Invest in training and tailored technical assistance** to foster AI literacy among human services providers. Training and technical assistance, such as conducting AI use case evaluations and promoting agencies’ responsible AI usage, can strengthen agencies’ AI skills and capacity. As a research and data analytics organization and technical assistance provider, Mathematica’s approach to using AI recognizes that the promise of AI comes with questions and concerns about data quality, ethics, security, confidentiality, and bias. Before using AI with or on behalf of our clients, we take the time to listen and understand. We want to know how they think about and apply AI, any concerns or questions they might have about the technology, their AI policies, and whether using the technology is suitable based on their specific program needs, goals, and the existing data.

3. How is AI already augmenting human efforts in the outlined transformational areas (or any other human service area you believe is ripe for transformation with AI)?

When deployed successfully, AI solutions support human decision making by streamlining processes such as data collection, synthesis of evidence bases, and summarization of written text. Importantly, these solutions do not replace expert-level judgment; they help bring information and evidence to the forefront and allow experts to focus their time more productively. Here, we aim to highlight several transformative areas that are particularly timely and salient in our work with health and human services agencies, practitioners, and participants:

- **Triaging tools for literature searches and evidence reviews.** AI-backed techniques can be used to develop triaging tools to support time-consuming research syntheses such as literature searches and evidence reviews. For example, Mathematica recently used an AI application to [support human experts during a literature review](#) assessing the broader impacts of scientific research. Under this approach, an expert reviewer curates an initial set of articles and determines whether each is relevant or irrelevant to the research question. The AI tool then uses these initial determinations to predict whether future articles are also likely to be relevant and sorts new articles in order of that likely relevance. As the expert reviewer continues to engage with the tool, each of their screening decisions feeds back into the predictive algorithm, and the AI tool dynamically updates the predicted relevance of each future article and re-ranks them. This approach holds great promise to perform literature searches and evidence reviews more efficiently and cost-effectively, allowing expert reviewers to focus their time more precisely on the evidence most likely to be relevant to the question at hand.

- **Early warning systems to target support to those in need.** AI models can be used to strengthen early warning systems and can be especially impactful in use cases that protect the rights and promote the welfare of the most vulnerable. Mathematica recently partnered with a state child welfare agency to gain insights into [which information available at first hotline communication is most associated with the need for a department safety response](#). Using predictive modeling techniques, we helped the agency identify which types of calls are associated with the most risk of impending danger to the child, and conversely, which are associated with child safety. This knowledge, in conjunction with the perspective of people who have lived experience with the child protection system, can inform policy on which cases do or do not merit investigation and help channel resources toward those in need.
- **Powering comprehensive policy scans.** AI methods can synthesize natural language, identify patterns, and unlock insights that might not be readily apparent to human reviewers. These approaches are particularly applicable to policy scans where relevant information might be spread across a range of websites or regulatory documents, making it difficult to get a clear picture of how policies are being implemented and enforced. Mathematica recently used AI techniques to [summarize the anti-harassment policies of institutes of higher learning](#) and measure the extent to which these institutions were adhering to agency guidelines. As AI technologies, especially large language models, grow increasingly performant, there is great potential to reduce the human effort to collect and codify large amounts of policy and regulation; in turn, enabling human services agencies, policymakers, and researchers to focus on critical questions, such as how disparate policy regulations or enforcement affect the experiences and outcomes of participants in human services programs.

4. Share recommendations you have for leveraging AI to advance human services mission outcomes.

The [Blueprint for an AI Bill of Rights](#) and the [AI Risk Management Framework](#) play key roles in codifying guiding principles for responsible AI, including addressing issues such as bias. These and other agency resources from the U.S. Department of Health and Human Services (HHS), such as its [Plan for Promoting Responsible Use of Artificial Intelligence in Automated and Algorithmic Systems by State, Local, Tribal, and Territorial Governments in the Administration of Public Benefits](#) and ACF's Policy on Generative Artificial Intelligence Tools (dated June 17, 2024) make up a growing ecosystem of relevant and timely AI resources to empower human services agencies, practitioners, and researchers seeking to implement AI solutions to address human service mission outcomes.

One way to build on these resources is to develop tailored practitioner's guides with built-in checkpoints. Efforts to integrate AI into human services and health agency operations, support for frontline staff, and research activities should include a well-defined framework for evaluating use cases. AI practitioners seeking to implement AI solutions should be prompted to provide a sound explanation for why AI adds value to their work and outline its advantages relative to other approaches. Tailored guidance based on predefined use cases or risk profiles could also outline key checkpoints to assess security, privacy, ethical, and equity considerations. Practitioner's guides might also provide clear guidance on agency coordination and technologies necessary to implement AI solutions.

We also recommend exploration of additional potentially transformative domains relevant to ACF's mission and work. Based on our experience collaborating with human services agencies, practitioners, and participants, we put forward these domains which could benefit from AI integration:

- **Enhance data linking and cross-agency interoperability.** Administrative data sources may exist in data siloes, without interconnections to other related data sources. This often poses challenges in addressing research questions and the key priorities of human services agencies. For example, agency staff might wish to link administrative data with case notes data to answer questions about the incidence rates for child maltreatment, as well as associated risk and protective factors. AI systems can better enable these types of links between structured administrative data and unstructured data sources, such as case notes recorded in natural language. ACF's [Interoperability Action Plan](#) provides a framework for data linking and interoperability initiatives; AI technologies hold promise to expand and accelerate these efforts toward cross-cutting research and evidence-based decision making accounting for the integrated system of human services programming and delivery.
- **Support review of competitive grant applications.** Notices of funding opportunities from ACF often generate large numbers of grant applications that cannot all be fully funded. Reviewing these applications might require coordination and collaboration between agency staff, experts in the field, and other key audiences impacted by the grant funding. AI systems could introduce efficiencies into this process. For example, an AI tool might be trained to flag applications missing key information, alerting applicants to potential issues in submissions in real time, and offering them an opportunity to make revisions. This process could ensure reviewers receive complete applications with all information necessary to make funding decisions.
- **Streamline resource publication processes.** Human services agencies publish resources regularly to disseminate research findings and guidance. These resources must be accessible, which in practice, involves dedicating resources to ensure documents are properly structured and content is properly tagged to be compatible with screen reader devices. New advances in AI such as [auto-tagging](#) hold great promise to promote efficiency while ensuring accessibility.

5. How should ACF and/or HHS incentivize innovative use of AI with appropriate safeguards?

AI innovation is a multifaceted process that must provide proper safeguards for equity, governance, privacy, and security. Frameworks and considerations to foster AI innovation are documented in the [White House Office of Management and Budget Memorandum M-24-10](#); the [HHS Plan for Promoting Responsible Use of Artificial Intelligence in Automated and Algorithmic Systems by State, Local, Tribal, and Territorial Governments in the Administration of Public Benefits](#); and the ACF Policy on Generative Artificial Intelligence Tools (dated June 17, 2024). The following ideas to incentivize innovation, influenced by our experiences with AI technologies and health and human services providers, build on the tenets of these frameworks:

- **Fund AI demonstration projects and evaluate their impacts.** Testing AI tools in a real-world context can help human services agencies gather information about equity, accuracy, and usefulness before deciding whether to deploy these tools on a larger scale. For example, ACF might fund demonstration projects for state and local TANF agencies to pilot innovative approaches to providing occupational supports for TANF recipients and other people with low

income. Establishing partnerships with researchers and technical assistance providers can strengthen the implementation of demonstration projects and ensure rigorous, reliable evaluations that are data-driven and incorporate feedback from program participants and community members. For example, researchers might assess equity implications by reviewing how an AI tool's predictive accuracy varies across demographic groups.

- **Hold a competition to incubate innovation and learn from practitioners in other industries.** Prize competitions have long been recognized as a mechanism to spur innovation, and have seen [increasing popularity at the federal level](#) in recent years. Competitions have the added benefit of [increasing the diversity](#) of voices and methodologies brought forward to tackle a specific problem, often attracting academic or industry competitors who might not interact regularly with the competition hosting agency. ACF could consider hosting a competition to solve an operational or research challenge. For example, ACF might make available a deidentified data set of Head Start participants and ask competitors to predict long-term educational outcomes, such as high school graduation. Modeling techniques that prove successful could be piloted as part of an early warning system to focus support on Head Start participants or sites most in need. Depending on the sensitivity level of the data in question, ACF might also consider advanced data-simulation techniques that can produce synthetic data with statistical properties sampled from real-world data. For example, Mathematica designed the competition data sets for the 2022 American Causal Inference Conference challenge [to mirror data from evaluations of large-scale U.S. health care system interventions](#) that aim to lower Medicare expenditures.

Mathematica Inc.

Our employee-owners work nationwide and around the world.

Find us at mathematica.org and edi-global.com.



Mathematica, Progress Together, and the “spotlight M” logo are registered trademarks of Mathematica Inc.