ATARC

Guiding Principles for Trustworthy Use of Generative AI:

Establishing Best Practices

Table of Contents

Government AI Stakeholders

Roles:

CIO, CTO, CFO, CISO, CDO, CAIO, OGC

Responsibilities:

- Key decision-making authority for AI initiatives and policies
- Ensure alignment with overall agency strategy Support agile decision-making with a small, focused board
- Empowered to veto decisions if necessary, similar to a UN Security Council mode

Roles:

Representatives, specialized in mission-specific AI applications, Data SMEs

Responsibilities:

- Provide expertise on niche or advanced AI use cases based on the agency's mission
- Collaborate with other board members on cutting-edge AI research and innovation
- Ensure compliance with AI standards and regulations

Roles:

Public Advocate (representing external users/ customers), Internal Advocate (representing agency employees), Representatives from regions, public, and user groups

Responsibilities:

- Ensure AI Systems meet the needs of end users (both public and internal)
- Promote transparency and accessibility of AI-driven services

Government AI Stakeholders

Roles:

Security, Network, Applications Teams, Data Scientists and Data Engineers

Responsibilities:

- Oversee AI technology infrastructure and data governance
- Ensure cybersecurity and network readiness for AI systems
- Drive AI innovations from an IT perspective

Roles:

Directors, Supervisors, Senior Managers, OIG/Auditor, Authorizing Official (AO, Risk Management Executive)

Domains:

Policy, Privacy, Finance, Civil Rights, Human Resources, General Counsel, Communications

Responsibilities:

- Expertise-driven oversight of AI-related matters in specific knowledge domains
- Address mission-critical AI use cases
- Advocate for the responsible and equitable use of AI

Generative AI Risk Assessment

ATARC

THREATS:

Risks to AI

Risks to System or Model Effectiveness

EXAMPLES

Data Poisoning:

The malicious alteration of data to make a system produce an adverse output

Cybersecurity:

Cyber threats to gain access or control to a system or exfiltration of system outputs

Lack of Infrastructure/ Storage Capacity:

Cyber threats to gain access or control to a system or exfiltration of system outputs

ATO Structure:

Slow authorization frameworks that require small tweaks to model to go through the entire authorization process again

Workforce Training:

A workforce without domain expertise or technical knowledge that can reduce model effectiveness or result in the wrong people being the humans-in-the-loop

TECHNICAL MITIGATIONS

Supply Chain Analysis:

Perform analyses of the system's supply chain including training data and require a bill of materials (BOM) stating where data originates

Continuous Monitoring Tools:

Use continuous monitoring tools to detect unauthorized access to models

Infrastructure Upgrades:

Use cloud-based infrastructure that can scale based on need easier than on-prem systems

Modular Systems:

Use modular architectures and pre-approved components that don't require constant reauthorization

Technical Experience:

Run trainings that focus on hands-on use of AI systems and simplified lessons on how the models work

GOVERNANCE MITIGATIONS

Frequent Data Audits:

Regularly audit the data of AI systems to ensure no major changes have been made without approval and that source data is clean

Zero Trust Architecture:

Use NIST 800-53 standards, SBOMs (Software Bill of Materials), and vulnerability management protocols to ensure data protection

Adaptable Cost Plan:

Infrastructure costs can vary depending on model usage, create a flexible plan to allocate budget for compute costs during times it's needed most

Reciprocity & Sponsorship:

Sponsor trusted partners for approval through authorities like FedRAMP, and build out reciprocity systems to streamline model approvals

Responsible AI Training

Train the workforce on the ethical dilemmas posed by AI, and arm them with the knowledge to intervene when necessary

Generative AI Risk Assessment

ATARC

BEHAVIOR:

Risks from AI

First and Second Order Risks Caused by the Model

EXAMPLES

Bias:	Misinformation/ Disinformation:	Intellectual Property:	Sustainability & Environmental:	Overreliance:	Rights-Impacting Decisions:	Safety-Impacting Decisions:
Statistical bias and discriminatory bias	Incorrect information distributed by systems either accidentally or maliciously--can also take the form of deepfakes	Questions about who owns synthetically generated content or use of copyrighted information a model is trained on	High-energy costs and impact on the environment	Relying on a system to make every decision with little human review--taking a model at its word for everything with little oversight	Consent, Notification, Privacy, and Civil Liberties	Threats to human health or wellbeing due to decisions or reccomendations made by a system

TECHNICAL MITIGATIONS

Bias Detection:	Watermarking:	Supply Chain Analysis:	Automated Metering:	System Warnings:	Technical Limits:	Automated Overrides:
Use tools that detect bias in AI outputs and run tests on data to eliminate statistical bias	Develop a watermarking system to identify synthetic content or deploy a content provenance standard	Perform analyses of the system's supply chain including training data and require a bill of materials (BOM) stating where data originates and if training consent was provided	Establish automated energy consumption cutoffs for on-prem systems to prevent high overrun costs	Include proper notification to users within the system itself about potential risks posed by the system	Ensure that user consent is given when interacting with a system. Use automatic notifications systems to tell individuals when decisions have been made by an AI system. Automatically delete user data that is not needed and ensure that data cannot be personally identifiable.	Establish automated overrides in systems to prevent it from taking certain courses of action without human approval

GOVERNANCE MITIGATIONS

Human-in-the-Loop (HITL):	Content Moderation:	Intellectual Property Guidelines:	Sustainability Standards:	Human-in-the-Loop (HITL):	Human-in-the-Loop (HITL):	Human-in-the-Loop (HITL):
Ensure humans are involved in decision-making done by AI in order to remediate bias concerns before a final decision is made	Ensure human oversight exists to both fact-check and prevent the spread of incorrect information produced by models	Establish clear guidelines for who owns content and data produced by AI in the contracting process	Institute contract requirements for sustainability that can include options such as carbon footprint tracking and carbon offsets	Ensure humans are involved in decision-making done by AI in order to remediate bias concerns before a final decision is made	Ensure humans are involved in decision-making done by AI in order to remediate bias concerns before a final decision is made	Ensure humans are involved in decision-making done by AI in order to remediate bias concerns before a final decision is made

Generative AI Risk Assessment

ATARC

BUSINESS:

Risks from Not Using AI

How Agencies Fall Behind without AI

EXAMPLES

Underreliance:

Not understanding AI and its uses could result in lost opportunities

Decreased Efficiency:

Repetitive tasks that could be automated pile up resulting in lower efficiency

Workforce Burnout:

Workers are stuck doing tasks that could easily be automated, resulting in burnout and talent loss

Lack of Innovation:

Unique approaches to challenges are never discovered, hindering innovation across an agency

Worse Decisions:

Decisions are not supported by data, leading to worse outcomes

TECHNICAL MITIGATIONS

Small-Scale Pilots:

Deploy small scale AI pilots and tackle low-hanging fruit first to better understand the benefits and limitations of AI systems

Adopt AI for Rote Projects:

Adopt AI models that can help automate repetitive processes such as document processing, contract writing, data analysis, etc.

Upskilling:

Adopt AI models that can help automate repetitive processes so that workers can be freed up for meaningful work they are passionate about

Use Generative AI for Permutations:

Adopt generative AI systems and use them as part of the brainstorming process to arrive at new and unique conclusions

Use AI for Data-Driven Insights:

Adopt AI systems to get access to relevant data quickly and make decisions with more accuracy and at greater speed

GOVERNANCE MITIGATIONS

Use Case Inventories:

Create AI Use Case Inventories to learn more about how other agencies are using AI

Incentives:

Create AI Use Case Inventories to learn more about how other agencies are using AI

Incentives:

Create AI Use Case Inventories to learn more about how other agencies are using AI

Generative AI Policy:

Institute a generative AI policy that encourages workers to ask questions to the model and avoid overregulation of the system that requires jumping through multiple hoops to access

Encourage AI Use by Decision-Makers:

Ensure that agency decision-makers leverage AI in the decision-making process and ensure they understand the benefits and limitations of the technology

ATARC

Managing Risks in Federal Procurement

Security Compliance

Security Assurance

Contracts should ensure that the proper cybersecurity requirements are in place for AI systems. These security measures should align with Zero Trust principles.

FedRAMP

FedRAMP is a rigorous security process that will help address the security of federal AI systems. AI solutions should achieve FedRAMP accreditation, however, many agencies have cited this process as a barrier to AI adoption in their compliance plans. Agencies should help sponsor unique solutions for FedRAMP approval to speed up this key process.

Incident Reporting

Contracts should require agencies and/or vendors to report a cybersecurity breach within 72 hours of the known incident.

Bill of Materials

Contracts should encourage vendors to supply a "bill of materials" that can trace the data provenance to prevent risks such as bias or data poisoning.

Security Overreach

Agencies should be careful not to overprescribe security requirements for low-risk systems that may hinder innovation.

Privacy

Privacy from the Beginning

Privacy measures should be included from the very start of the solicitation process. Agencies should include their Senior Agency Official for Privacy (SAOP) and/or the agency privacy program during the pre-solicitation phase.

Anonymization & Deletion

Agencies should request information from vendors regarding their measures in place to anonymize personally identifiable information (PII), delete temporary data, and ensure models cannot produce outputs that would leak sensitive PII.

Data Transfer

Agencies should help protect users by ensuring vendors cannot transfer or share data outside the specified purpose of the model

Managing Risks in Federal Procurement

ATARC

Performance Monitoring

Testing Transparency

Contracts should ensure that vendors are transparent with training, validation and testing data and that both vendor and agency are working off the same information.

Agency Audits

Agencies should have free access to conduct audits of models and testing data to ensure systems are working as advertised.

Real-World Context

All testing should be conducted in conditions as close to real-world as possible to better understand system performance.

Intellectual Property Protection

Synthetic Content Generation

Agencies should encourage vendors to provide information about the copyright status of data which models have been trained on to avoid copyright infringement in system outputs.

Data & Model Ownership

Agencies should clearly delineate who owns both the training data and the model in contracts to avoid being locked out of their own data or system by contractors.

Unauthorized Data Training

Agencies should ensure that their data cannot by used by vendors without their express consent for use in systems other than the one specified in the contract.

Competition in Procurement

Transparency of Models

Agencies should use open-source models where possible to promote interoperability and sharing of information.

Data and Model Portability

Agencies should take measures to avoid vendor lock-in of their data and systems.

Pricing Transparency

Contract terms should include price transparency and allow for the sharing of these pricing terms with other agencies.

Notification

Notification Mechanism

Contracts should include a mechanism that identifies the proper channels individuals can be notified if they are impacted by an AI-enabled decision or issue with PII or civil liberties. It should also outline which party is responsible for notification.

User-Friendly Notification

Notification mechanisms should be user-friendly and intuitive. Agencies should avoid long terms of service or confusing language. Systems should notify users when they're using AI and should be prioritized in interactions where humans may believe they're interacting with another human and not a machine.

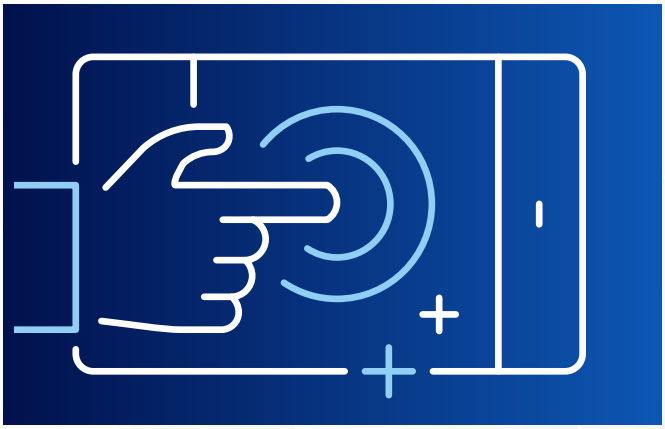

Human-in-the-Loop (HITL)

Considerations

Throughout AI lifecycle for oversight, especially in high-risk use cases *(e.g., decision-making, policy writing)*

To mitigate bias or fix issues in outcomes

Ensure human review in critical decision points *(e.g., legal or regulatory matters)*

The human must have the knowledge and experience to both choose the correct model for the task and know when intervention is necessary

Continuous monitoring capabilities need to be put in place to adequately notify and provide the right information for a human intervention

Hypothetical Scenario: SNAP Application Processing

Mary has recently lost her job and has turned to SNAP benefits to help get her back on her feet. The SNAP program has recently adopted an AI-enabled system to help process applications. When Mary enters her documents into the system, it identifies her as ineligible for SNAP benefits due to an income discrepancy. The AI system misinterprets her unemployment benefits as full-time income and denies her application.

Human-in-the-Loop Intervention

01

Initial Flagging for Human Review

A human reviews the denied application and the explanation given by the AI system.

02

Understanding the Issue

The human identifies the document that has been misinterpreted by the AI system that it based its denial on. The human realizes the system misinterpreted the document.

03

Fixing the Mistake

The human manually updates Mary's application based on the correct information, allowing it to be approved.

04

Preventing Further Risk

The human notifies the model developers of the issue so that it can be adjusted to prevent a similar scenario from occurring in the future.

Outcome

Without human intervention Mary's application would have remained denied, and she would likely face financial hardship and place additional burdens on the program through the appeals process. Human intervention allowed Mary to get her benefits quickly, avoid any legal challenges, and improve the existing model.

Choosing RAG and Fine-Tuning

Considerations

Retrieval-Augmented Generation (RAG): Suitable for applications needing real-time updates from other sources, ideal for dynamic scenarios like customer inquiries. RAG allows you to use a large language model to query trusted, targeted data. RAG also can provide the user to accurately cite the source of information provided in outputs

Fine-Tuning: Recommended when deep knowledge is essential, such as specialized legal or technical domains requiring context-driven outputs

Cost: Variations of RAG (e.g., using iteration) exist that may present cost/efficiency tradeoffs that must be evaluated

Hypothetical Scenario: FEMA Disaster Response

A Category 5 hurricane makes landfall, causing widespread destruction across multiple states. FEMA's AI system is tasked with helping responders prioritize resource allocation, provide situation reports, and answer queries from local officials.

Retrieval-Augmented Generation (RAG)

FEMA's AI system uses RAG to access relevant and trusted data from various government databases that provide information such as weather, infrastructure damage, and local resources.

A responder is looking to direct injured victims on the ground to the hospital that is both close and has capacity for the victims. The responder asks the AI system "Which hospitals in the affected area have power and bed capacity?" The AI system gathers the latest information from FEMA's database and provides an answer based on the data. The model provides an answer that can be quickly acted on since it was trained on FEMA disaster response guidelines and procedures.

Non-RAG Fine-Tuning

FEMA's AI system has also been fine-tuned on agency-specific heuristics, rules, and frameworks.

A responder is looking to identify the proper legal policies and the right channel to deploy resources to an affected area. The responder asks the system "Which disaster assistance programs can be activated for this affected community?" The AI system which has been fine-tuned to FEMA policies identifies that this specific community is eligible for Individual Assistance, Public Assistance, and Hazard Mitigation Assistance. The system uses FEMA terminology and policies to draft up a press release and memorandum.

Outcome

RAG allows access to up-to-date information from trusted sources, while the fine-tuning of the model displays this information in a manner that aligns with FEMA's policy framework. By using both methods, FEMA is able to guide first responders based on recent information and is able to activate the proper resources based on FEMA policy.

Without RAG, responders would not have access to the right data to make a decision due to the rapidly changing situation on the ground. Without fine-tuning, responders would have to align outputs to FEMA policy manually and adjust terminology to fit their standards, creating excess administrative work and delays in providing assistance.

Categorizing Use Cases by Risk

Considerations

Rights Impacting AI: AI whose output serves as a principal basis for a decision or action concerning a specific individual or entity that has a legal, material, binding, or similarly significant effect on that individual's or entity's civil rights, civil liberties, privacy, equal opportunities, or access to government services

Safety Impacting AI: AI whose output produces an action or serves as a principal basis for a decision that has the potential to significantly impact the safety of human life or wellbeing, climate or environment, critical infrastructure, or strategic assets or resources

Non-Rights and Safety Impacting: AI models that do not fall into the above categories are considered "Non-Rights and Safety Impacting".

Hypothetical Scenario: Department of State Workflow Efficiency

The Department of State is looking to streamline a number of its processes using artificial intelligence. To do so, the agency has adopted two new solutions: a generative AI solution for Freedom of Information Act (FOIA) requests and a generative AI solution to help translate technical language in policies and documents into more plain language. Each of these solutions will require different resources and attention based on the risks associated with them.

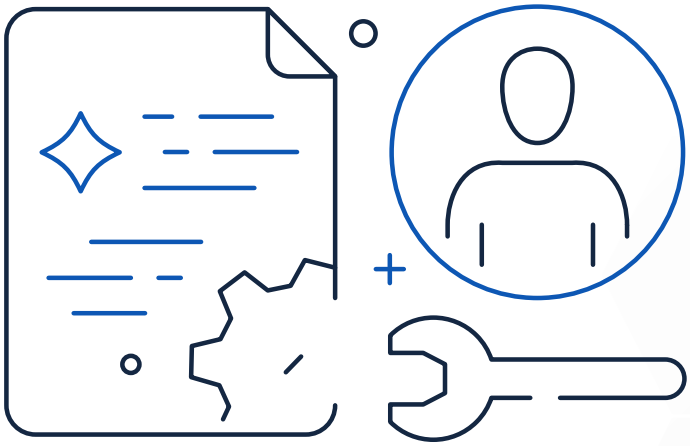

Safety and Rights-Impacting AI Systems

The FOIA processing solution assists in reviewing requests, collecting relevant documents, and suggesting redactions to human users. This system had to pass rigorous testing and requires significant continued oversight to ensure safe operation. However, this system could impact safety and rights since it might miss sensitive information or inadvertently disclose private personally identifiable information (PII). It also may miss documents that are responsive to the FOIA request. Therefore, it will probably need more oversight to ensure safe operation than a system simplifying document language.

Non-Safety and Rights-Impacting AI Systems

The plain language solution helps users by taking highly technical documents or policies and summarizing them in language anyone can understand. This system is a fairly innocuous use of artificial intelligence and is unlikely to have wide-ranging safety or rights implications. Minimal oversight is required due to low risk and impact if the summarization is not for rights-impacting documents or policies.

Outcome

Classifying AI systems into low risk and high risk categories allows for a more efficient allocation of resources such as knowledgeable humans-in-the-loop for oversight and testing. While the high-risk systems receive special deployment pathways and enhanced monitoring, the low-risk systems can be deployed more rapidly.

Without this separation, the low-risk systems would likely face overregulation slowing down their time to deployment, while the high-risk systems would not receive enough oversight and dedicated resources.

Standardizing Best Practices

Considerations

Reinforcement Learning: Use checklists (e.g., Section508.gov) as a foundation for automating

Permutations: Use LLMs to explore unconsidered solutions

Building on Existing Practices: Leverage existing data governance, privacy, and security frameworks to manage AI risks

Mitigating Bias: Clearly distinguish between statistical bias (algorithmic) and social/economic bias and implement governance frameworks to mitigate both types of bias

Use Case Solutions Library: Centralize successful AI use cases to guide future innovation

Hypothetical Scenario: AI in Veterans Healthcare

The VA implements an enterprise-wide AI initiative to improve veterans' healthcare and institutes a number of best practices to help it achieve success.

Best Practices in Action

01

Reinforcement Learning:
Automating Repetitive Tasks Using Checklists

The VA helps automate important tasks like accessibility standards using a reinforcement learning algorithm. The algorithm is trained on standards from Section508.gov and use a checklist derived from these standards to identify issues with accessibility design. This algorithm fits the purpose and helps free up time for staff to tackle more important issues.

02

Permutations:
Exploring Unconsidered Solutions with LLMs

The VA is looking to optimize the allocation of doctors to rural clinics that are facing staffing shortages. They implement a large language model (LLM) to help provide innovative solutions to the issue. The LLM suggests combining telemedicine with rotating on-site visits, an initiative that the department did not previously explore due to assumptions about logistical issues.

03

Building on Existing Practices:
Leveraging Data Governance Frameworks

The VA wants to ensure that its AI systems follow strict privacy and security standards. The agency looks to existing data governance frameworks like the Federal Information Security Management Act (FISMA) and Health Insurance Portability and Accountability Act (HIPAA) to help guide their management of AI systems.

04

Mitigating Bias:
Addressing Statistical and Social/Economic Bias

The VA wants to use an AI system to predict the risk of a veteran developing a chronic condition based on historical data. The agency identifies two different types of bias: ***statistical bias*** (ex: underrepresentation of certain groups in the dataset) and ***social/economic bias*** (systemic disparities in the healthcare system). The agency implements governance frameworks for both of these biases to help mitigate possible harm caused by an AI system trained on these datasets.

05

Use Case Solutions Library:
Centralizing Successful AI Use Cases

The VA is looking to replicate its successful AI implementation across the agency. They add detailed information on each use case to their AI use case inventory, including a description, cost, benefits, POCs, and open source code to allow others in the agency to replicate the success for their own use cases.

Outcome

By following best practices and sharing information not only within the VA but learning from other agencies' examples, the VA successfully implements it's AI initiative. If the VA did not implement these best practices, their systems would likely be siloed, innovation would be stifled, and public trust would be undermined.

Additional Resources

Human-in-the-Loop (HITL)

Alexia, Gaudeul., Ottla, Arrigoni., Vicky, Charisi., Marina, Escobar-Planas., Isabelle, Hupont. (2024). Understanding the Impact of Human Oversight on Discriminatory Outcomes in AI-Supported Decision-Making. Frontiers in artificial intelligence and applications, doi: 10.3233/faia240598

Marion, Ho-Dac., Baptiste, Martinez. (2024). Human Oversight of Artificial Intelligence and Technical Standardisation. arXiv.org, abs/2407.17481, doi: 10.48550/arxiv.2407.17481

Susan, Fischer. (2024). 3. Questioning AI: Promoting Decision-Making Autonomy Through Reflection. doi: 10.48550/arxiv.2409.10250

Sunday, Abayomi, Joseph., Titilayo, Modupe, Kolade., Onyinye, Obioha, Val., Olubukola, Omolara, Adebiyi., Olumide, Samuel, Ogungbemi., Oluwaseun, Oladeji, Olaniyi. (2024). 4. AI-Powered Information Governance: Balancing Automation and Human Oversight for Optimal Organization Productivity. Asian Journal of Research in Computer Science, doi: 10.9734/ajrcos/2024/v17i10513

Choosing RAG and Fine-Tuning

K., Rangan., Yiqiao, Yin. (2024). 1. A Fine-tuning Enhanced RAG System with Quantized Influence Measure as AI Judge. doi: 10.48550/arxiv.2402.17081

Maria, A., de, Luis, Balaguer., Vinamra, Benara., Renato, Luiz, de, Freitas, Cunha., Roberto, de, M., Estevao, Filho., Todd, Hendry., Daniel, Holstein., Jennifer, Marsman., Nick, Mecklenburg., Sara, Malvar., Leonardo, Nunes., Rafael, Padilha., Morris, Sharp., B., Silva., Swati, Sharma., Vijay, Aski., Ranveer, Chandra. (2024). 2. RAG vs Fine-tuning: Pipelines, Tradeoffs, and a Case Study on Agriculture. arXiv.org, doi: 10.48550/arxiv.2401.08406

Róbert, Lakatos., P., Pollner., András, Hajdu., Tamas, Joo. (2024). 3. Investigating the performance of Retrieval-Augmented Generation and fine-tuning for the development of AI-driven knowledge-based systems. arXiv.org, doi: 10.48550/arxiv.2403.09727

Harshit, Kumar, Chaubey., Gaurav, Tripathi., Rajnish, Ranjan., Srinivasa, k., Gopalaiyengar. (2024). 4. Comparative Analysis of RAG, Fine-Tuning, and Prompt Engineering in Chatbot Development. doi: 10.1109/icftss61109.2024.10691338

Categorizing Use Cases by Risk

Yi, Zhang., Kevin, Klyman., A, Zhou., Youngjae, Yu., Min, Pan., Ruoxi, Jia., Dunlun, Song., Percy, Liang., Bo, Li. (2024). 1. AI Risk Categorization Decoded (AIR 2024): From Government Regulations to Corporate Policies. doi: 10.48550/arxiv.2406.17864

Eli, Sherman., Ian, W., Eisenberg. (2023). 2. AI Risk Profiles: A Standards Proposal for Pre-Deployment AI Risk Disclosures. arXiv.org, doi: 10.48550/arxiv.2309.13176

Standardizing Best Practices

Isha, Mishra., Vedika, Kashyap., Naresh, Kumar, Yadav., Rajesh, Pahwa. (2024). 1. Harmonizing Intelligence: A Holistic Approach to Bias Mitigation in Artificial Intelligence (AI). Deleted Journal, doi: 10.47392/irjaeh.2024.0270

Kin-Ho, Lam., Delyar, Tabatabai., Jed, H., Irvine., Donald, B., Bertucci., Anita, Ruangrotsakun., Minsuk, Kahng., Alan, Fern. (2022). 2. Beyond Value: CHECKLIST for Testing Inferences in Planning-Based RL. doi: 10.48550/arXiv.2206.02039

ATARC would like to take this opportunity to recognize the following Generative AI Working Group members for their contributions:

Frank Indiviglio • *National Oceanic and Atmospheric Administration, ATARC Generative AI Working Group Government Chair*

Nathan Manzotti • *Federal Deposit Insurance Corporation, ATARC Generative AI Working Group Government Vice Chair*

KJ Lian • *Amazon Web Services, ATARC Generative AI Working Group Industry Chair*

Michael Adams • *ATARC Generative AI Working Group Carahsoft Chair*

Anthony Boese • *ATARC AI and Data Policy Working Group Government Chair*

Joshua Iseler • *Carahsoft*

Dr. Tanya Kuza • *Veterans Affairs*

Chris Robinson • *General Services Administration*

John Sprague • *National Aeronautics and Space Administration*

Dartagnan Fischer • *Department of Homeland Security*

John K. Wright II • *National Oceanic and Atmospheric Administration*

Nicholas Rappold • *National Weather Service*

Natalie Buda Smith • *Library of Congress*

Brian H. Seborg • *University of Maryland Baltimore County Emeritus*

Henry Sienkiewicz • *Georgetown University*

Ken Farber • *Teksynap*

Dr. Marcus Weller • *AIONIC Machine Learning*

Sandy Barsky • *Oracle*

Amanda Henninger • *The Fraclex Group*

Jordan Jannone • *California State University, Northridge*

Greg Crabb • *Ballistic Ventures*

Patrick Stingley • *ATARC Generative AI Working Group Member*

Keith L. Moore • *Martin-Blanck & Associates*

Brian Pfeifer • *Department of State*

David Egts • *Salesforce*

Wendy Marquez • *Wize Solutions*

Mun-Wai Hon • *Northern Virginia Community College Department Chair*

Norman Wong • *Palo Alto Networks*

Trisha Christian • *Small Business Administration*

Brian Peretti • *Treasury*

Akhtar Zaman • *National Archives and Records Administration*

Dr. Ferdous Khan • *Department of Homeland Security*

MG Karch • *ATARC Generative AI Working Group Member*

About ATARC

ATARC serves as a dynamic bridge connecting government, industry, and academia to collaboratively address emerging technology challenges. By fostering open dialogue and introducing cutting-edge solutions, ATARC equips Federal agencies with essential market research opportunities and actionable insights to navigate the evolving technology landscape. At its core, ATARC is dedicated to driving technological innovation and building partnerships that transcend traditional boundaries. Through impactful events, in-depth white papers, forward-looking research initiatives, and collaborative Working Groups, ATARC creates an ecosystem for knowledge sharing and problem-solving.

About The Generative AI Working Group

The mission of ATARC's Generative AI Working group is to leverage government, industry, and academic experience with the procurement and deployment of Generative AI capabilities to mitigate business risk, improve employee productivity, and enhance the quality of customer experience when consuming government services.The Generative AI working group aspires to promote modernization and innovation by sharing modern practices of implementation, exploring new technologies, and providing guidance to achieve IT transformation success through Generative AI.

ATARC