Memorandum

From: **USC-Amazon Center for Secure and Trusted Machine Learning**
To: **USC Viterbi Faculty**
Subject: **Call for Proposals**
Date: **3/1/2024**

Description of Opportunity

USC and Amazon have created a joint research center focused on the development of new approaches to machine learning (ML) privacy, security, and trustworthiness. The Center for Secure and Trusted Machine Learning (in short, Trusted AI), which will be housed at the USC Viterbi School of Engineering, will support USC and Amazon researchers in the development of novel approaches to privacy-preserving ML solutions.

Each year, through a competitive process, the Center will provide support for several research projects focused on the development of new methodologies for secure and privacy-preserving machine learning solutions. The Center will also provide annual fellowships to doctoral students working in this research area, enabling them to advance research frontiers. Fellowship recipients will be named as Amazon ML Fellows in recognition of their promise and achievements. In addition to funded research projects and annual fellowships for doctoral students, the collaborators will host an annual joint public research symposium to share their knowledge with the machine learning and AI communities.

About Amazon

Amazon is improving customers' lives with practical, useful generative AI innovations. We do this by building and deploying AI across three technology layers: at the bottom layer we offer our own high performance and cost-effective custom chips, as well as a variety of other computing options including from third-parties. At the middle layer, we offer customers choice by providing the broadest selection of Foundation Models—both Amazon-built as well as those from other leading providers. At the top layer we offer generative AI applications and services to improve every customer experience.

There are three things that distinguish Amazon's approach to the development and deployment of AI:

1. Maintaining a strategic focus on improving the customer and employee experience through practical, real-world applications of AI.
2. Marshaling our world-class data, compute, and talent resources to drive AI innovation.
3. Committing to the development of responsible, reliable, and trustworthy AI.

We hereby invite USC faculty to submit proposals for the next round of funding via the USC-Amazon Center for Secure and Trusted Machine Learning.

Topics of Interest

Topics of interest would include, but are not limited to, those below. Please feel free to bring your/your institution's unique viewpoint and expertise to these topics:

Responsible Generative AI

- This may include, but is not limited to measurement and mitigation, guardrail models, privacy concerns, detecting and mitigating adversarial use cases, and machine unlearning and model disgorgement
- Responsible AI for audio, image and video generation
- Privacy preserving continual learning/self-learning
- Fact Checking and Factual Error Correction for Truthful LLMs

Large Language Models (LLMs):

- Retrieval augmented generation (RAG), fine-tuning and alignment (SFT, RLHF), and efficient inference: ensuring accuracy and reducing hallucinations; maintaining privacy and trust; reasoning over long contexts;
- Long form context methods
- Improving data efficiency; effectively distilling models for real-time inference, data quality checks
- Multi-lingual LLMs and challenges for cross-language defects (e.g. cross-language hallucinations)
- Synthetic data generation for LLM learning
- Adapting LLMs for dynamic content (e.g., feeds, web content) in knowledge-augmented scenarios
- Tool and Code Empowered LLM
- External Knowledge and Domain Knowledge Enhanced LLM and Knowledge Updating

Vision-Language:

- Multimodal learning and video understanding: retrieval with multimodal inputs (e.g., video, image, text, speech);
- Adversarial ML with multimodal inputs
- Comprehensive video understanding with diverse content (open-vocabulary).
- Shared multimodal representation spaces, aligned codecs
- LLM and VLM based Intelligent Agents

Search and Retrieval:

- Personalization in Search, semantic retrieval, conversational search: understanding descriptive and natural language queries for product search; retrieving information using LLMs' output
- Search page optimization (ranking) using heterogeneous content such as related keywords, shoppable images, videos, and ads
- Tool Learning for Proactive Information Seeking

Efficient Generative AI:

- Novel model architectures for improved performance (accuracy & efficiency)
- Training large neural network models with efficiency: High performance distributed training and inference algorithms for Generative AI systems, quality metrics and evaluations

Timeline

1. **Release of CFP:** 01 March 2024
2. **One page abstracts due**: 01 April 2024
3. **Full proposals due:** 24 May 2024 at 5pm PT
4. **Announcement of selected projects:** August 2024
5. **New projects start date:** August 2024

Scope, Funding, Eligibility, Evaluation Process, and Expectations

- Performance period will be one year, with the possibility of extending for another year.
- The per-project available funding is approximately $75k for sponsored gifts with no overhead.
- It is anticipated that about 4 projects will be selected for funding.
- All USC faculty who can serve as a PI on federally-funded projects are eligible.
- All proposals will be evaluated by an advisory committee composed of USC faculty and researchers from Amazon.
- The PIs of the selected projects are expected to
    a. provide quarterly reports as well as a final report of their research accomplishments in the project;
    b. participate and present at a technical USC-Amazon workshop on Trusted AI (planned for Fall 2024);
    c. participate and present at an annual Research Symposium on Trusted AI (planned for Spring 2025).

Abstract submission instructions

Interested USC faculty submit a **1 page maximum overview** (excluding references) of a potential full proposal. The purpose of these abstracts is for Amazon to review and provide feedback on the proposal: would Amazon be interested in collaborating with the researcher, could the proposal include new use cases for the science, and does the project align with industry goals. High level feedback will then be shared with each submitting faculty member. Feedback will help guide the faculty member towards submitting a more relevant proposal or a more impactful proposal (unless the abstract already draws strong interest from Amazon) when the full proposal deadline arrives on May 20. Abstracts are not a prerequisite for submitting proposals.

Abstracts in pdf format should be sent via email to the administrative program manager Ariana Perez (arianape@usc.edu), by 5 pm PT on 04/01/2024.

Proposal submission instructions

The full proposal should be prepared and submitted according to the following guidelines: format is single-spaced, 11-point font or larger, with no less than 0.5-inch page margins. The proposal should include the following content:

1. Project description (3 pages max), including the title, PI(s), an executive summary, technical description of the project, expected deliverables/outcomes, and milestones
2. List any USC background Intellectual Property (unlimited)
3. References (unlimited)
4. Requested budget, assuming a gift award with no overhead.
5. Bio(s) of the PI(s) (up to 2 pages per PI in NSF format)

Full proposals (compiled in a single-file PDF format) should be sent via email to the administrative program manager Ariana Perez (arianape@usc.edu), by 5 pm PT on 05/24/2024.