Amazon - Virginia Tech Initiative for Efficient and Robust Machine Learning Research Proposals
Frequently Asked Questions

Q: What is the Amazon - Virginia Initiative?
A: A partnership to advance research and innovation in artificial intelligence and machine learning. The Amazon - Virginia Tech Initiative for Efficient and Robust Machine Learning supports machine learning-focused research projects, doctoral student fellowships, community outreach, and an establishment of a shared advisory board.

Q: Where is the program located?
A: The initiative is housed in the College of Engineering and led by Sanghani Center for Artificial Intelligence and Data Analytics researchers on Virginia Tech's campuses in Blacksburg, Virginia and at the Innovation Campus, in Alexandria, Virginia.

Q: How can faculty get involved?
A: As part of the initiative, Virginia Tech faculty members will be invited to submit machine learning sponsored research projects for Amazon that work to revolutionize the way the world uses and understands this field of modern technology. Research project proposals may include support for students and staff. The initiative began accepting (and awarding) proposals from Virginia Tech faculty in the College of Engineering for research projects with support beginning Fall Semester 2024 (FY 2025).

Q: Who can submit a proposal?
A: The Amazon-Virginia Tech Initiative invites Virginia Tech faculty from the College of Engineering and faculty from the Department of Statistics and Department of Mathematics in the College of Science to submit proposals. Please refer to the Call for Proposals for more details.

Q: What are the research topics of interest?
A: The topics of interest for this year's round of proposal submissions include (but not limited to):

Context:
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; and
3. Committing to the development of responsible, reliable, and trustworthy AI.

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:

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

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

Q: Will proposals outside of the research topics of interest be considered?
A: Yes. While the above topics are particularly of interest, highly meritorious proposals in other areas of ML and AI will also be considered.

Q: How will proposals be selected?
A: Proposals will be evaluated by a collaborative advisory board composed of Virginia Tech faculty and Amazon researchers for their technical merits, potential to advance research in areas of ML, and opportunity for impact.

Q: What is the funding amount for the research projects?
A: Project funding is expected to be in the range of $50K - $100K. Proposals selected for funding will be awarded either as gifts or sponsored awards (will be decided based on project goals and objectives). The initiative reserves the right to negotiate awards to fulfill overall program goals.

Q: What does the project funding support?
A: Proposals may request funding for undergraduate students, graduate students, post-doctoral research associates, faculty, travel, equipment, and/or other expenses as deemed necessary for the successful execution of the project.

Q: Can faculty submit more than one proposal?
A: A faculty member can submit multiple proposals but can only serve as a PI on at most one proposal to this solicitation (and Co-PI on others).

Q: What is the proposal submission process?
A: The proposal submission process has two stages:
1. **Abstract:** Interested VT 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. Criteria at this stage include: interest in collaboration with the researcher, alignment with Amazon objectives, and the use cases the project can support. 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 us).
   a. Amazon plans to provide feedback on all abstracts. They will set up an office hours type format for PIs to breifly meet with abstract reviewers to provide feedback. Abstracts can not be rejected, and an abstract is not a prerequisite for submitting a proposal. The purpose of the abstracts is to give the PIs the opportunity to get feedback from Amazon scientists on the direction and relevance of their potential proposal.
2. **Full Proposal:** The full proposal (single-spaced, 12 pt font, 1" margins, 3 pages max) should be submitted as one PDF document comprising:
   i. Motivation,
   ii. Statement of the Problem,
   iii. Proposed Activities and Timeline,
   iv. Primary Team Members and Roles,
   v. Background IP (if none, state so), and
   vi. Simple ROM Itemized Budget with justification. A formal OSP budget is not needed. Include all direct costs but only direct costs (do not include overhead calculations in your budget).
   vii. References (1 page max; not included in page limit)
   viii. CVs of Investigators (2 pages max per investigator; not included in page limit)

Q: How do you submit the proposal?
A: Proposal submissions and any questions about proposal submissions should be addressed to Naren Ramkarishnan, Director of the Amazon-VT initiative at amazon-vt@cs.vt.edu.

If you have any additional questions about the Amazon - Virginia Tech Initiative for Efficient and Robust Machine Learning or the proposal submission process, please reach out to Naren Ramkarishnan, Director of the Amazon-VT Initiative at naren@cs.vt.edu.