

Amazon - Virginia Tech Initiative for Efficient and Robust Machine Learning
Call for Research Proposals

In 2022 Virginia Tech and Amazon announced their 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.

The Amazon-Virginia Tech Initiative for Efficient and Robust Machine Learning (ML) is soliciting proposals from Virginia Tech faculty in the College of Engineering for research projects with anticipated support beginning Fall Semester 2024 (FY 2025). The objective of this initiative is to advance research and innovation in Artificial Intelligence (AI) and ML, specifically in the areas of efficient and robust machine learning. Virginia Tech faculty members are invited to submit proposals for research projects that work to revolutionize the way ML is developed, evaluated, and embedded in everyday contexts, devices, and applications (see below for list of topics of interest).

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

While the above topics are particularly of interest, highly meritorious nominations in other areas of ML and AI will also be considered.

Proposals should be submitted by a Virginia Tech faculty member serving as a PI and can include 1-2 Co-PIs. Proposals may request support for undergraduate students, graduate students, postdoctoral research associates, faculty, travel, equipment, and/or other expenses as deemed necessary for the successful execution of the project.

Projects cannot exceed one year in duration. Projects with longer-running aims will need to apply again for future year(s) of funding. Proposed budgets are expected to be in the range of \$50 - 100K direct costs (with most awards expected to be in the middle of this range). 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).

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. 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.

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 our Amazon partners 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).
 - a. We plan to provide feedback on all submitted abstracts. We will set up an “office hours” type feedback format for PIs to briefly meet with abstract reviewers to receive feedback. Abstracts are not accepted or rejected, and an abstract is not a prerequisite for submitting a proposal. However, submitting an abstract gives PIs the opportunity to obtain direct 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)

Timeline

- March 18, 2024 (5pm EDT): 1 Page Abstract Due
- April 2, 2024 Information Session open to the VT community (Amazon will aim to have all abstract feedback delivered ahead of the April 2nd info session, but this is subject to change)
- April 26, 2024 (5pm EDT): Full Applications Due
- June 26, 2024: Decisions Announced
- Aug 10, 2024: Projects begin in the Fall Semester
- July 31, 2025: PIs submit a 2-page report

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.

Review Criteria

A joint VT-Amazon advisory board will review applications based upon

- Novelty
- Relevance to Alexa
- Practicality

If you have any additional questions about the Amazon - Virginia Tech Initiative for Efficient and Robust Machine Learning or the proposal submission process, please refer to the following FAQ.