

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 2025 (FY 2026).

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:

- **Foundation Model Improvements**
 - Novel model architectures
 - Novel training algorithms and methodologies
 - Multi-modal (e.g., text, image, video, audio) understanding and generation

- Multi-lingual understanding and generation
- Algorithms and workflows for acquiring and curating high-quality and diverse datasets for training
- **Foundation Model Evaluation**
 - Creation of new benchmarks for assessing foundation model capabilities
 - Methodologies for robust evaluation of generative AI systems, including agents
- **Efficient Generative AI**
 - Efficient training and inference for cloud and on-device applications
 - Compute and memory efficient handling of multi-modal long/infinite context
 - Improving efficiency of diffusion models, including discrete diffusion
- **Reasoning**
 - Commonsense and domain-specific (e.g., math, coding) reasoning
 - Temporal and spatial reasoning
 - Reasoning for planning
- **Knowledge Grounding**
 - Approaches to ground (multi-modal) generation on up-to-date world, domain-specific, enterprise, or personal knowledge
 - Memory-augmented generative AI systems
- **Agentic AI**
 - Creation of autonomous systems capable of performing tasks, making decisions, and interacting with their environments
 - Multi-Agent systems and agent orchestration framework improvements
 - Customization and continual improvement of agents post deployment
 - Internationalizing agentic systems
- **Responsible Generative AI**
 - Red teaming (e.g., advanced red teaming approaches, automated red teaming for multi-modal models)
 - Improvement of foundation model RAI performance (e.g., robustness to jailbreaking and membership inference attacks; watermarking approaches; deepfake detection)
 - Responsible agentic AI (e.g., robustness of multi-agent systems, adherence to guardrails)
- **Applications of Generative AI**
 - Systems that leverage Generative AI for advancing science and technology in areas such as physics, mathematics, chemistry, biology, hardware design, materials science, engineering, economics, healthcare, climate.

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: Projects will typically be funded at a range of \$50 - \$100K direct costs, depending on the project scope and nature of funding. If a project is pursued as Sponsored Research, the [Amazon-VT Initiative] Advisory Board will discuss this adjustment in conjunction with the proposer, and ensure that applicable indirect costs are covered for awarded projects. Sponsored Research Project recipients will not be asked to reduce their budgets to accommodate indirect costs.

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 briefly 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 abstract?

A: Proposal abstracts should be submitted via this Google Form link: <https://forms.gle/dS4nbY38TRdnXuGo9/>. Please fill out all the questions in the form before attaching your abstract.

Q: How do you submit the proposal?

A: Full proposals should be submitted via this Google Form link: <https://forms.gle/MC33f2h7H5VEFYFd8>. Please fill out all the questions in the form before attaching your proposal. Any questions about abstract/proposal submissions should be addressed to Wanawsha Shalaby, Program Manager 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 abstract/proposal submission process, email Wanawsha Shalaby, Program Manager of the Amazon-VT Initiative at amazon-vt@cs.vt.edu.