

Amazon - Virginia Tech Initiative for Efficient and Robust Machine Learning Ph.D. Fellowship
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 students get involved?

A: The new partnership will create opportunities for Virginia Tech graduate and doctoral level students who are interested in and currently pursuing educational and research experiences in artificial intelligence-focused fields. Named Amazon Machine Learning Fellows in recognition of their scholarly achievements.

Q: Who qualifies for the Amazon ML Fellowship?

A: Students must be enrolled in a PhD program at Virginia Tech, be in good standing, and should have exhibited outstanding academic performance to be eligible for a fellowship. Additionally, students must be in the third, or fourth year of their Ph.D. studies and pursuing doctoral-level research in machine learning, data science, AI, and/or NLP.

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: How can students apply for the Amazon ML Fellowship?

A: Students must be nominated for the fellowship by Virginia Tech faculty members. Nomination materials are to be submitted by the Virginia Tech faculty member. Refer to the Call for Proposal for more details about the application process. ****Students cannot submit their own nomination package****

Q: What funding is provided to students?

A: Amazon ML Fellows receive a \$24K stipend for one academic year (two semesters), full tuition coverage, and travel support of \$1K to attend conferences in student's area of research.

Q: What additional benefits do Amazon ML Fellows receive?

A: Each fellow will be invited to interview for a paid Amazon internship during the summer after the fellowship year. If selected, awardees will gain valuable industry insight and experience through direct engagement with Amazon researchers.

Q: How do you submit the nomination?

A: Nominations should be sent in one file via PDF 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 Amazon ML Fellowship, please reach out to Naren Ramkarishnan at naren@cs.vt.edu.