

**CapitalOne-Illinois Center for Generative AI Safety, Knowledge Systems, and
Cybersecurity (ASKS)
Call for Proposals 2025-2026**

Abstracts Due: May 9, 2025
Full Proposals Due: June 13, 2025

1. INTRODUCTION

Generative Artificial Intelligence Models (GAIMs) have emerged as transformative tools, driving advancements across diverse domains. However, their widespread adoption raises critical concerns regarding trustworthiness across dimensions such as truthfulness, safety, cybersecurity, fairness, robustness, and privacy. These models have demonstrated remarkable performance on knowledge reasoning tasks, owing to their implicit knowledge derived from extensive pretraining data. However, their inherent knowledge bases often suffer from disorganization and illusion, bias towards common entities, and rapid obsolescence. Consequently, GAIMs frequently make up untruthful information, exhibit resistance to updating outdated knowledge, or struggle with generalizing across multiple languages. The Capital One-Illinois Center for Generative AI Safety, Knowledge Systems, and Cybersecurity (ASKS) is a collaboration between Capital One and University of Illinois Urbana-Champaign, aims to advance GAIMs' up-to-date knowledge, safety, truthfulness, trustworthiness across multiple dimensions and model types, and paving the way for safer and more responsible integration of GAIMs into critical applications including cybersecurity and situation understanding.

2. CFP 2025-2026 TOPICS

The ASKS Center is calling for proposals for the second funding round, 2025-2026 academic year. In this phase, the center will support 3-5 sponsored projects. Each funded project will support a PhD research assistant and 2-4 weeks of summer salary for principal investigators (PIs). Each selected project will then engage with research scientists from Capital One to collaborate and jointly advise PhD students.

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

Large Language Models (LLMs) and Vision Language Models (VLMs) Safety:

- LLM Reasoning, System 2 Thinking
- LLM agents
- Creative Intelligence
- Guardrail models and agents
- Synthetic Data Generation
- Jailbreaking, prompt injection, and their defenses
- Inference attacks, extraction attacks, and their defenses
- Hallucination Mitigation, LLM and VLM factuality
- Multimodal knowledge representation
- Knowledge localization, updating and editing

- Debiasing and Unlearning
- Theme-based and structure-enhanced RAG (retrieval augmented generation)
- Privacy Preservation
- Human-Model Alignment

Knowledge Systems:

- Open-domain Never-Ending Knowledge Extraction, Discovery and Acquisition
- Knowledge base construction and operation
- Anomaly Event Detection
- Situation Understanding, Forecasting and Reporting
- LLM-enhanced information extraction and knowledge graph construction

Cybersecurity:

- Multi-dimensional anomaly detection
- LLM for Explanation
- Time Series Data Mining
- Temporal Spatial Aware Foundation Models
- AI for cybersecurity
- Privacy-preserving data analytics

3. PROPOSAL GUIDELINES

- Eligibility: Full-time tenure-track, research-track and teaching faculty members at Illinois are eligible to submit proposals as PIs.
- Submissions should not reference either Capital One, nor speculate about how the research might be applied to current or future Capital One products, services, business models or needs. Proposals should focus on science.
- Stage 1 abstract submissions are due **May 9, 2025**. Interested Capital One faculty submit a **1-page maximum overview** (excluding references) of a potential full proposal. The purpose of these abstracts is for Capital One to review and provide feedback on the proposal: would Capital One be interested in collaborating with this researcher, could the proposal include new use cases for the science, and does the project align with industry goals. Feedback will help guide the faculty member towards submitting a more relevant and impactful proposal. Abstracts are not a prerequisite for submitting proposals. A PDF of the 1-page abstract should be submitted to: [Link](#) on or before **May 9, 2025, at midnight CDT**.
- An ASKS abstract feedback session will be held between **May 12-May16, 2025**. The Capital One Science review team will give feedback on whether a full proposal is encouraged for each abstract submission. It is hoped that this exchange will serve as a matchmaking exercise, enabling the development of well-focused proposals that are aligned with the interests of the PI and Capital One.
- Stage 2 full proposal submissions are due **June 13, 2025**. PIs who did not submit abstracts are still eligible to submit full proposals. The proposal should not exceed three pages, with unlimited references. PIs are encouraged to seek collaborations with research scientists at Capital One to prepare for proposals. Proposals will be evaluated by a collaborative advisory board composed of ASKS Leadership Team (excluding board members who submitted proposals) and Capital One scientists for their technical merits and innovations,

topic relevance, potential to advance research in focus areas, and broader impact.

- Full Proposal Format: The proposal 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:
 - Full names and email addresses of all PIs involved
 - Project description (3 pages max), including the focus area of proposal (per Call for Proposal process) title, PI(s), an executive summary, technical description of the project, expected deliverables/outcomes, milestones and what plans exist for open sourcing data or results
 - List all university background IP (unlimited)
 - References (unlimited)
 - Requested budget (1 page max), use format shown in Section 7
 - Biographies of the PIs (up to 3 pages per PI in NSF Format)
- Selection criteria: Successful projects will be evaluated on the promise and progress in the quality of publications, student mentoring, and potential technology transfer activities.
- Project period of performance: September 1, 2025–August 31, 2026.

4. IMPORTANT DATES

- April 18, 2025: CFP Topics Announced
- May 9, 2025: Abstract Submission Deadline
- May 9 – May 31, 2025: Abstract feedback sessions with Capital One scientists and UIUC Faculty
- June 13, 2025: Full Proposal Submission Due
- Early July 2025: Full Proposal Acceptance Notification
- TBD September 2025 (On Campus ASKS Center Event)

5. SUBMISSION AND CONTACT INFORMATION

PDFs of the submissions documents should be uploaded to [Here](#). For questions about the proposal submission process, please contact [asks-management@lists.cs.illinois.edu]. For questions about the ASKS Center, please contact the center director Prof. Heng Ji [hengji@illinois.edu] or the leadership team: [asks-leadership@lists.cs.illinois.edu].

Commented [HJ1]: Setup link

6. ASKS LEADERSHIP TEAM

- Director: Heng Ji (Illinois)
- Center Liaison: Prem Natarajan (Capital One)
- Program Manager: Katie Clark (Capital One)
- Center Alliance Manager: Akshay Mulera (Illinois)
- Advisory board members: Jiawei Han (Illinois), Bo Li (Illinois), Ling Ren (Illinois),

Commented [HJ2]: Add capitalone advisory board members

7. BUDGET FORMAT

Work with your department business office to complete a budget table in the format shown below. The categories and text in the “basis of estimate” column are provided as examples; please update to reflect actual expenses and project needs.

Per the Collaboration Agreement, no F&A costs are to be assessed on Gift-Funded Research Projects. If, through the review process, the project is deemed to be a Sponsored Research Project, budgets will be renegotiated.

Item	Amount	Basis of Estimate
PI Summer Salary	\$XX	2 weeks summer salary, including fringe benefits
GRA Salary Support	\$XX	X-months GRA at Y% effort, including fringe
Other Direct Costs	\$XX	List items. E.g., hard drives, high-speed cluster access, publication fees, software costs, travel
Total	\$XX	