NATIONAL SCIENCE FOUNDATION
2415 EISENHOWER AVENUE
ALEXANDRIA, VIRGINIA 22314

NSF 23-105

Frequently Asked Questions (FAQs) for Program Solicitation NSF 23-562, for Safe Learning-Enabled Systems

1. **What is the exact limit on the number of proposals per person?**

   For each deadline and category of this solicitation, an investigator is allowed to participate in only one (1) proposal, serving as either the Principal Investigator (PI), co-Principal Investigator (co-PI), Project Director, Senior Personnel, or Consultant. Additionally, throughout the duration of the program (FY 2023 - FY 2024), an investigator is not permitted to be PI, co-PI, Project Director, Senior Personnel, or Consultant on more than two (2) awards.

2. **What is a sub-scale system?**

A sub-scale system is a smaller version of the larger system that is designed to replicate some or all of the functionality of the larger system. For example, in the aerospace industry, sub-scale systems are commonly used to test new aircraft designs before building a full-scale prototype. These sub-scale systems can include wind tunnel models, miniature propulsion systems, or scaled-down cockpit simulators.

3. **The notion of safety can include a variety of different aspects. What is the scope of this program?**

   For the purposes of this program, safe learning-enabled systems are broadly construed as ones that do not exhibit catastrophic behaviors during their execution. A proposal must define its notion of safety and must justify why this notion is important. Note that notions such as trust, perceived safety, security, privacy, fairness, accountability, ethics, are outside the scope of the program.

4. **Does the program favor any specific applications?**

   The focus of this program is on safety concerns related to learning-enabled systems. It is not the program's intention to attribute safety problems to any particular applications or systems. However, proposals must clearly articulate the notion of end-to-end safety that is either mathematically- or empirically-based using clear and unambiguous language.

5. **How many awards will be made in each category?**

   We have an estimated budget of $10,000,000 for each fiscal year dedicated to this program. The number of awards in each project class will be determined based on the quality of proposals received within that category. Please note that no funds have been specifically allocated to either project class at this point. Given the breadth of research topics included in this program, we anticipate a highly competitive competition.

6. **What is an end-to-end system and end-to-end safety?**

   An end-to-end system is a complete and integrated system that encompasses all stages of a process, from the initial input to the final output of its operating environment. End-to-end safety refers to safety properties defined for an end-to-end system. For example, a system may be built out of several components, or in layers. End-to-end means “across all components/layers”.

7. **What is an example of "unknown unknown"?**

   The term "unknown unknown" refers to a type of knowledge or information that is not only unknown but also not recognized as unknown. For example, a research team has developed a new system and tested it extensively but fails to anticipate the impact of a certain environmental event that was not considered during testing, possibly because the event is rare. That event would be an unknown unknown.

8. **Does the program favor any particular model (system), e.g., ChatGPT, or exclude any machine learning methods, e.g., Adversarial ML? Are only static learning systems in scope, or must we require continual learning?**

No. This program supports fundamental research to address safety problems in learning enabled systems. It neither focuses on any particular model or system, nor excludes any machine learning methods as long as the proposal addresses meaningful safety problems as described in the solicitation. The proposal must provide a precise, unambiguous definition of safety and motivate the definition in the context of the application under consideration. See the solicitation, particularly Program Description (Section II) and Proposal Preparation and Submission Instructions (Section V).

9. **How are the proposals reviewed in this program?**

   The proposals submitted to this program will be reviewed by NSF based on NSF Merit Review policy contained in Chapter III of the *NSF Proposal and Award Policies and Procedures Guide* (PAPPG). Additional detailed information on the review process can be found in the solicitation. Open Philanthropy and Good Ventures may send their representatives to the review panels as observers, but these representatives cannot serve as panelists or reviewers.

10. **Will Open Philanthropy and Good Ventures provide funding or be involved in funding decisions?**

    Open Philanthropy and Good Ventures will provide funding for this program. However, funding decisions will be made by NSF, based on NSF processes and NSF program officer recommendations.

11. **Will end-to-end security strategies for a deep learning system be considered for this solicitation?**

    The program is about safety of learning-enabled systems and not about security of such systems, hence research on security is not in scope. To the extent that a deep learning system may cause unsafe behavior by acting upon insecure data, the proposal would be in scope.

12. **Are methods that improve our ability to inspect safety issues of AI methods but not necessarily guarantee safety in scope?**

    No.

13. **Can the proposed work leverage existing systems already built by the members of the team?**

    Yes.

14. **Would improving a learned system's error on out of distribution data be in scope?**

    Yes. The proposal must motivate/justify how reducing the error would improve safety.