

Google Research

Research Scholar Program
2025 Application Companion

tl;dr – This Application Companion provides resources to support applicants with structuring and developing high-quality applications. It includes the program overview, how to apply and formatting guidelines, samples of high-scoring responses, open advice to proposers and a submission FAQ. It is available at goo.gle/rs-application-companion.

For more information, visit our [website](#) or contact research-awards@google.com.

Table of Contents

Program overview

- Award Details

- Review criteria

- Eligibility criteria

How to apply

- Locating the application

- Formatting guidelines

FAQ

Sample responses to current application questions

- Proposal Abstract

 - Is your research consistent with Google’s AI Principles? Why or why not?

 - How are you currently working to broaden participation of individuals from groups traditionally underrepresented in computing, please provide specific examples?

Open advice to proposal writers

Program overview

The Research Scholar Program champions groundbreaking discoveries by empowering the next generation of leading academics. Through unrestricted gifts, the program provides vital support for research initiatives at esteemed institutions worldwide.

This funding is specifically dedicated to fostering the innovative work of early-career professors, recognizing their potential to make significant contributions to their respective fields. By investing in these rising stars, the Research Scholar Program seeks to accelerate the pace of knowledge creation and address some of the world's most pressing challenges

Research Scholar is committed to addressing evolving global challenges through research. This year, we invite you to submit proposals in the following areas:

Algorithms, Optimization, and Markets : Algorithms and optimization form the foundations of computer science, focusing on designing efficient methods to solve complex contemporary problems including problems with applications in machine learning, data science, and modern AI. The primary goals in this area are to create methods that improve resource efficiency and sometimes offer guarantees on the quality of the solution. Key

focus areas include combinatorial optimization, market algorithms, operations research, continuous optimization and learning, Scalable algorithms, and general topics in theoretical computer science. This line of research is crucial since it studies the solvability of problems through a set of tools that nicely complement machine learning techniques.

For our area, we call for proposals specifically in the areas of:

- Combinatorial optimization
- Market algorithms
- Operations research
- Continuous optimization and learning
- Scalable algorithms
- Other

Applied Science: Large language, visual, and multimodal models have made significant advances in recent years, opening up new possibilities for scientific research. The program will evaluate proposals based on their scientific merit, creativity, and potential impact on the field of scientific research.

For our area, we call for proposals specifically in the areas of:

- Applications: Proposals that demonstrate how large language models can be used to advance scientific discovery in a specific field.
- Foundations and agentic tasks: Proposals that explore broad advances in building, tuning, or deploying large models for scientific research, such as integrating language models with specialized scientific tools, developing agent based models that address complex multi-step tasks potentially requiring code generation, developments of models that help in synthesis of scientific literature, and accelerating scientific analysis, experimentation, and summarization with or without humans in the loop.
- Evaluation: Proposals that develop datasets or methods for benchmarking and evaluating large models for science, including evaluation of coding ability using appropriate libraries, evaluating domain-specific knowledge, assessing factuality and grounding, evaluating multimodal capabilities, and developing tasks that require multi-step scientific reasoning.
- HCI: Proposals that enhance scientific workflows, such as automating complex simulation pipelines, with large language models and human-in-the-loop interaction.

Health research: Google's Health research aims to advance AI and technology that helps people live healthier lives. Achieving this goal will require collaborative research with public officials, clinicians, and consumers. In partnership with public officials, we are creating tools to understand population level health. With clinicians, we are developing novel algorithms to better understand and make use of complex medical data such as images, text, lab tests, and genomics. With consumers, we are developing technology that helps people find high quality health information and better understand their own health status. By focusing on inclusive, transformative research we aim to improve the lives of billions of people.

For our area, we call for proposals specifically in the areas of:

- Generating and understanding large datasets of the world to derive useful insights for improving population health, especially in under resourced regions or communities

- Novel algorithm development for better understanding of complex medical data, with focus areas in novel methods, novel applications, or underserved settings
- Novel methods, including both software and hardware, that helps extract health insights cheaper, faster, or better

Human-computer interaction: HCI researchers at Google design and build large-scale interactive systems that aim to be humane, simple-to-understand, and delightful to use. We work across a variety of HCI disciplines, including predictive and intelligent user interfaces, mobile and ubiquitous computing, social and collaborative computing, and interactive visualization.

Machine Perception: Machine perception researchers at Google develop algorithms and systems to tackle a wide range of tasks, including action recognition, object recognition and detection, hand-writing recognition, audio understanding, perceptual similarity measures, and image and video compression.

Machine learning: Machine learning, a cornerstone of Google's research initiatives, encompasses a vast spectrum of exploration. This includes fundamental theoretical investigations into algorithms and their underlying principles, as well as the development of practical applications that address real-world challenges. Through these diverse research endeavors, Google aims to advance the state-of-the-art in machine learning and harness its potential to drive innovation across a wide range of domains.

For our area, we call for proposals specifically in the areas of:

- Learning algorithms & techniques
- Learning theory
- Federated learning
- Information theory
- Optimization for ML algorithms
- Reinforcement learning
- Robotics
- Recommender systems

Natural language processing: Our team comprises multiple research groups working on a wide range of natural language understanding and generation projects. Our researchers are focused on advancing the state of the art in natural language technologies and accelerating adoption everywhere for the benefit of the user. Natural Language Processing and understanding plays a major role in driving Google's company-wide OKRs as Language understanding is the key to unlocking Google's approach: "Build a more helpful Google for everyone that increases the world's knowledge, success, health, and happiness."

Networking: Networking is central to modern computing, from WANs connecting cell phones to massive data stores, to the data-center interconnects that deliver seamless storage and fine-grained distributed computing. Because our distributed computing infrastructure is a key differentiator for the company, Google has long focused on building network infrastructure to support our scale, availability, and performance needs, and to apply our expertise and infrastructure to solve similar problems for Cloud customers. Our research combines building and deploying novel networking systems at unprecedented scale, with recent work focusing on fundamental questions around data center architecture, cloud virtual networking, and wide-area network interconnects. We

helped pioneer the use of Software Defined Networking, the application of ML to networking, and the development of large-scale management infrastructure including telemetry systems. We are also addressing congestion control and bandwidth management, capacity planning, and designing networks to meet traffic demands. We build cross-layer systems to ensure high network availability and reliability. By publishing our findings at premier research venues, we continue to engage both academic and industrial partners to further the state of the art in networked systems.

For our area, we call for proposals specifically in the areas of:

- Cloud Networking: We constantly evolve cloud networking solutions to provide a great cloud experience to billions of users. Our focus area covers customer-facing networking API design to the network data and control plane programming including HW programming. We exercise the Hybrid Research model by deploying our solutions in the Google Cloud Platform, which is one of the largest and fastest-growing cloud providers in industry. Our activities include networking, distributed systems, network security, kernel programming and algorithms.
- WAN networking (a.k.a. GGN) is responsible for the design, development, build and operation of Google's global network that every Google service runs on. GGN develops cutting-edge networking technologies that allow Google's global WAN to be zero touch, builds out some of the largest scale Software Defined Networks (SDNs) infrastructure ever deployed, develops sophisticated software systems for network capacity forecasting, planning and optimization, designs and implements new optical technologies. GGN relies on the most advanced techniques in network hardware and software, traffic engineering, and network management to deliver unprecedented scale, availability and performance at industry leading cost points. Additionally, we are also advancing the state of the art in data analytics and machine learning to drive network efficiency and optimization at scale.
- Network Infrastructure for Data Centers brings together networking, distributed systems, kernel and systems programming, end-host stacks, and advanced algorithms to create the datacenter networks that power Google. We deploy real-world systems at a global scale.

Privacy, Safety, and Security: Google Privacy, Safety, and Security is committed to ensuring that the internet is safer for everyone. To meet our goal, we support and partner with academia to bring about state of the art advancements across a broad range of privacy, security, and safety areas.

For our area, we call for proposals specifically on:

- Novel applications of AI for privacy, security, and safety
- Ensuring the privacy, security, and safety of AI systems
- User and measurement studies of privacy, security, and safety
- Applied cryptography
- Differential privacy
- Hardware security and side-channel analysis
- Software vulnerabilities, software supply chains, and fuzzing

Topics outside of these areas will still be considered. However, we encourage applicants to align their proposals with one of the above topics.

Quantum Computing: Two primary goals of the Quantum AI team are to develop a functional quantum computer that can tolerate errors and to identify novel applications that can be executed on quantum hardware. We actively collaborate with academic partners to advance these goals and we welcome the submission of proposals containing innovative ideas.

For our area, we call for proposals specifically in the areas of:

- Quantum algorithms
- Quantum machine learning
- Quantum error correction
- Early fault-tolerant quantum computing
- NISQ experiments
- Superconducting qubits
- Neutral atom quantum computing

Software engineering and programming languages: Research on all aspects of software development, including the engineers and the programming languages, libraries, development tools, and processes that they use. This includes software development methodologies and tools, debugging practices and tools, software testing strategies and tools, cooperation strategies for developers, interface and library design, code optimization and verification techniques, etc.

Systems: Google's systems and networking systems research is focused on building and deploying novel systems at unprecedented scale. Our work spans the entire spectrum of computing, from large-scale distributed systems to individual machines to accelerator technologies.

We address fundamental questions around data center architecture, cloud virtual networking, wide-area network interconnects, software-defined networking, machine learning for networking, large-scale management infrastructure, congestion control, bandwidth management, capacity planning, and designing networks to meet traffic demands.

Award Details

We encourage submissions from professors globally who are teaching at universities and meet the eligibility requirements. It is our hope that this program will help develop collaborations with new professors and encourage the formation of long-term relationships. Awards are disbursed as unrestricted gifts to the university and are not intended for overhead or indirect costs. They are intended for use during the academic year in which the award is provided to support the professor's research efforts.

Review criteria

Only complete applications that meet the following criteria will be scored:

- Submitted by eligible applicants
- Related to computing or technology in one or more focus areas
- Adhere to the required formatting guidelines

Eligibility criteria

- The PI and/or co-PI must be full-time professors (associate, assistant, etc.) at degree-granting institutions.
- The applicants must have received their PhD within seven years of submission.
- A faculty can apply a maximum of three times within the seven years.
- Faculty can qualify for funds through this program only one time and they can be a previous FRA recipient.

Scoring will be based on the following areas:

- **Faculty Merit.** Track record of publishing at top tier conferences, engagement in the research community, potential to execute research projects successfully, contributions to the open source community.
- **Research Merit.** Research work aligns with Google research and is an area of interest to Google, it advances state of the art and is good for the world, innovative from prior work in the field, probability of impacting our products/engineering/service
- **Proposal Quality.** The proposal is clear, focused and follows guidelines. Committee found it easy to navigate, answers the question of how they plan to approach the problem and their capability of executing the research successfully, provides further documentation of potential impact from the program
- **AI Ethics Principles.** The research proposal supports Google's AI Principles (e.g. the social benefit is clear and/or supports scientific excellence in the relevant field) and there are no outstanding concerns related to the other Principles.
- **Broadening Participation:** The proposal demonstrates a strong commitment to broadening participation in computing through initiatives such as mentoring underrepresented groups, establishing collaborations with diverse institutions, and disseminating research in accessible formats.

How to apply

Locating the application

Between December 20th 2024 and January 27th, 2025, applicants can apply by clicking the "Apply now" button on the program website, which will direct applicants to Submittable, where they will need to login or create a new account to begin an application.

Please be sure to complete all required fields, upload any necessary documents, and thoroughly review your application before clicking "Submit."

Formatting guidelines

Filename: Submit your proposal and CV as a single PDF file. Name the file in the following format: "[First InitialLast name]-2025" (e.g., "JDoe-2025"). Use only letters, numbers, and hyphens.

Proposal length:

- Sole PI: Maximum 3 pages for the proposal, 2 pages for the CV (total of 5 pages).
- Co-PI: Maximum 3 pages for the proposal, 2 pages for each CV (total of 7 pages).

Formatting:

- Single-spaced
- 1-inch margins
- Times New Roman 12-point font

Proposal should include the following numbered sections:

Overview (3 pages max)

1. Proposal Title
2. Full name, contact information (postal address, email address, phone), and affiliation (institution and/or department) of PI(s)
3. Abstract (concise summary of proposal)
4. Research goals and problem statement
5. Description of the proposed work, expected outcomes, and results
6. Discussion of how the research relates to prior work (including your own)
7. Explanation of your qualifications to conduct this research
8. For ongoing projects, explain how this funding would enhance your existing project.
9. Data policy: Describe your intentions for sharing the project's output with the broader research community (e.g., open-sourcing code, making datasets public). Please note that for those awards that are structured as unrestricted gifts, there are no legal requirements once a project is selected for funding. This is simply a statement of your current intentions. However, for research area topics that are not awarded as unrestricted gifts (usually those that require the use of a specific product, methodology, or other constraint), open sourcing the software, models, or other intellectual property developed during the project will be a mandatory condition for receiving the award, unless otherwise specified in a separate agreement between Google and the recipient.

CVs (4 pages max):

1. Primary PI: 2-page max CV required
2. Co-PI: 2-page max CV (optional)

Important notes

1. The co-PI's CV is the only content allowed on the additional 2 pages of a co-PI proposal. Any submitted CV longer than 2 pages will be truncated before review.
2. Proposals without a co-PI's CV should not exceed 5 pages.
3. References should be excluded from the proposal itself. Instead, use the designated sections in the application form for this information.

Sample responses to current application questions

Research Scholar Program is launching once again this year, and high-scoring proposals from this cycle will be made available in future cycles. For your reference, we've included high-scoring samples from previous cycles that should offer guidance in crafting your own compelling Research Scholar proposal.

Proposal Abstract

Sample 1: Large Language Model (LLM)-based technologies, such as voice assistants and chatbots, have become commonplace tools for work, school, entertainment, information and resource retrieval, and personal assistance. To minimize harm during interaction with these tools, their underlying training mechanisms (i.e., training data and annotation) must accurately represent diverse users, particularly those already societally marginalized. This proposal advocates for reflexive and collaborative practices in data annotation between data annotators and members of marginalized communities to identify and mitigate bias at multiple stages of LLM development. In this work, we focus specifically on the disability community and plan to (1) generate a dataset of realistic disability-related conversations between people with disabilities (PWD) and an LLM-based chatbot and (2) investigate how data annotators label this dataset before and after a reflexive discussion and joint annotation session with PWD. This project will contribute a dataset of chatbot conversations with PWD and a framework for joint annotation processes with marginalized communities.

Sample 2: With the growing trend of outsourcing machine learning (ML) computation in security-critical and sensitive domains, there is a pressing need to provide accountability to ML model supply chains. This project proposes a novel conceptual approach to authenticate ML models for use in cases where (i) users remotely querying a model require assurance that they are accessing the intended model, and (ii) auditors must hold model providers accountable for deploying a flawed model. Our research will comprehensively assess the suitability of potential technical solutions for developing an accountable ML model deployment. Furthermore, we will design and evaluate novel accountability mechanisms that can be implemented through collaborative efforts between external auditors and ML model providers. The resulting solution will represent the first realistic response to the existing lack of accountability in the ML pipeline.

Sample 3: The high operational costs present a significant challenge for major technology companies looking to profit from their LLM-based intelligent programming assistants (LLM-IPAs). For example, GitHub Copilot, with approximately 1.5 million users, reportedly incurs an average loss of more than \$20 per user per month, even when users are charged \$10 per month for the service subscription according to a report from WSJ. The high operational cost also comes along with high carbon emissions. Therefore, reducing the cost of LLMs has

become imperative for the sustainability of the industry and our environment. This project proposes novel and effective solutions to speed up LLM inference by developing a more compact, LLM-oriented program representation and designing a specialized dynamic inference mechanism for LLM-assisted programming. It reduces the number of tokens that need to be processed or generated by LLMs and avoids redundant model computation with finer-grained inference management, making LLM-assisted programming faster, greener, and cheaper.

Is your research consistent with Google's AI Principles? Why or why not?

Sample 1: This work is consistent with Google's AI Principles. The project centers on social good, fighting against harm and bias in AI datasets, and involving community members in research to reinforce accountability for positive user representation in the algorithms we hope to influence. This research aims to identify and mitigate harm in AI tools, poses minimal risk to users, is unrelated to weaponry and surveillance, and does not violate any human rights principles. This research has just been submitted to the Institutional Review Board for Wellesley college via Brandeis University and will likely fall under exempt or expedited review.

Sample 2: Yes, the project is consistent with Google's AI principles. A key objective of the project is to develop technologies that would improve accountability of deployed ML models, which aligns with Google's principles of accountability. Additionally, the project would have social benefits, particularly in mitigating bias and privacy risks within ML models. Regarding the ethics of the research, there are no concerns related to the proposed research; it does not involve human subjects and does not harm any of the stakeholders of the research by design. Therefore, the project has not undergone a review by an Ethics committee. The PI is fully aware of the privacy and security risks involved in scientific activities and is committed to ensuring the safety of the research while maintaining the highest standards of scientific excellence in the resulting research artifacts.

Sample 3: This research is consistent with Google's AI Principles. By improving the energy efficiency of large-language-model-based intelligent programming assistants, this project aims to optimize their computation resource efficiency and reduce carbon emissions, well aligned with the principle of being socially good. It also offers hints on developing more helpful prompts to fulfill developers' expectations, thus being accountable to people. The software and benchmarks developed throughout the project will be made open-sourced and available for uses that accord with these principles.

How are you currently working to broaden participation of individuals from groups traditionally underrepresented in computing, please provide specific examples?

Sample 1: -As an Assistant Professor at a historically women's college, and as a woman of color (WOC) myself, I dedicate my professional career to empowering young women in computer science (CS). I teach Wellesley's introductory CS course to a diverse group of women and strive to communicate that programming and CS research is within reach for each of them. I also teach Wellesley's Human-Computer Interaction course, a subject that broadens participation due to its approachable content.

- I support students outside the classroom by judging for the Wellesley Hackathon, serving on grad school advice panels, guiding individual applications for CS and information science programs, and participating in Q&A sessions with the CS Club. I also support senior theses of women in the CS department, primarily through co-advising and as a committee member.

- In my faculty research, I am collaborating with a mentor who is an incredible postdoctoral researcher at NYU, and the co-PI of this proposal. As I finish my first semester as a faculty member, I have four students interested in joining my upcoming lab, all of whom are WOC.

- My research employs participatory design methods with underrepresented communities to involve them in the technology development pipeline. I work closely with local disability organizations, education centers, and schools to learn from their expertise and give their members a voice in computing and design research venues (ACM ASSETS, CHI, and FAccT).

Sample 2: I regularly engage with students and faculty from institutions and organizations serving underrepresented groups. This year, I delivered a guest lecture in Spanish at INCA (INtroductory school in CryptogrAphy, colocated with Latincrypt, <https://incacrypto.org/2023>). The audience primarily consisted of K-12 and undergraduate students from Latin America. I also contribute to the MD4SG group, where I have served as a reviewer for conferences like MD4SG and EAAMO. Through the MD4SG working groups (<https://www.md4sg.com/workinggroups/fairness>), I've collaborated with and mentored PhD students from underrepresented backgrounds. In addition, in my role as a Lecturer at the University of Edinburgh, I regularly engage in mentoring and supervising PhD, master's, and undergraduate students from underrepresented groups. Finally, I regularly contribute as a volunteer by delivering introductory talks on operational security and applied cryptography to environmental activist organizations.

Sample 3: I am dedicated to supporting women in STEM. Among my research publications, I actively collaborate with female researchers, with at least one female co-author in over half of my published works. A notable achievement is our USENIX'23 paper, where we had more female co-authors (3) than male co-authors (2). This is a rarity in our community, especially considering that USENIX is among the most competitive security conferences. Moreover, I had the privilege of being invited as a panel member to the launch event of UniTy (formerly known as The Collective) at my Faculty, an initiative focused on supporting women in IT, to share my experience of working in STEM as a female researcher.

Open advice to proposal writers

As reviewers, we seek proposals that clearly articulate the problem and solution. Ensure your proposal demonstrates how your research in computing and technology aligns with our mission to create positive impact.

Consider the following when crafting your proposal:

1. Clearly define the problem. Good research starts with a compelling question.

2. Describe a specific, achievable outcome. What will this research enable that wouldn't happen otherwise, and how? Outline both minimum expected and best-case scenarios, specifying the datasets and test cases you'll use.
3. Differentiate your contribution. Clearly explain how your work advances the state of the art, using citations and other standard practices.
4. Outline your approach. Explain your plan for addressing the research challenges, even if all answers aren't yet known. Identify potential risks and mitigation strategies.
5. Contextualize the work. Describe existing funding and how this proposal fits into your broader research goals. How will this research be used? Will it build research capability, create a tool, reproduce a result, foster collaboration, follow up on an idea, or explore a new one? We are interested in all possibilities.
6. Make it accessible to non-experts. While we try to have your proposal reviewed by a Google expert in your field, it will also be read by non-experts, so please ensure the motivation and outcomes are understandable to a broad audience.

Your proposal should ultimately demonstrate how your research aligns with our mission to recognize and support academic researchers whose work in computing and technology makes a positive difference in the world.

FAQ

I need support with navigating Submittable. Where should I go?

Please visit the Submittable FAQ page at <https://www.submittable.com/help/submitter/> and [submit a ticket](#) at the bottom of the page if your question is not listed.

I need clarification on the Research Scholar Program's application guidelines or the questions themselves. Where should I go?

Please do not submit questions about the application or program through the Submittable message center as these inboxes are not monitored. Instead, try these steps:

1. Review the Research Scholar [FAQ](#) page located at the bottom of the Research Scholar page.
2. If your question remains, email research-awards@google.com and a member of our team will be glad to help.

How can I withdraw my application?

1. [Sign in](#) to your Submittable account.
2. Choose the Submissions tab in the main navigation bar or click your username in the upper right corner of the screen and select Submissions.
3. Click on the Active tab to see your active submissions. Click on the title of the submission that you wish

to withdraw.

4. When the submission opens, click the Withdraw link in the upper right area:
 - a. When the dialog box appears, type a Reason for withdrawal that will be received by the organization:
 - i. Click the Withdraw button. Your submission status will now be set to Withdrawn. Note: If your submission is in a terminal status (Accepted, Declined, or Completed) you will instead see a Remove link. Click it to remove the submission from your list view.