

Chan Zuckerberg Initiative

REQUEST FOR APPLICATIONS

Accelerating and Scaling Biological Sciences with AI

The Chan Zuckerberg Initiative invites applications to build large-scale AI/ML models using CZI's high-performance computing cluster to power new approaches to biological discovery.

OPPORTUNITY

The Chan Zuckerberg Initiative (CZI) invites proposals to build large-scale models that cannot be created with conventional university resources. Through competitive allocation of computing power on CZI's new GPU cluster, we will support inspired and cutting-edge model building that will power new approaches to biological discovery. Priority will be given to models that align with CZI's work to [build virtual cells](#) in partnership with the scientific community, but all proposals relating to [CZI's mission](#) to cure, prevent, or manage all diseases by the end of the century will be considered.

CZI's vision is to build the future of science by advancing biomedical research and leveraging advances in AI. Over the next 10 years, we're working to understand the mysteries of the cell and how cells interact within systems, which will underpin discoveries that will improve medicine in the decades that follow. CZI is the only philanthropic organization to fund and build one of the largest computing systems dedicated to nonprofit life sciences research in the world. CZI is making this cluster available to academic researchers across the U.S. through this RFA.

CZI's vision demands that progress in biology accelerates. Consequently, we are leveraging the potential of AI to build [virtual cell models](#): multi-scale, multi-modal, large neural network-based models that can represent and simulate the behavior of molecules, cells, and tissues across diverse states to help scientists explore the molecular underpinnings of human health and disease. We foresee virtual cells as having broad applications for biomedical research, disease diagnosis, and therapeutic development.

CZI values open science. Successful proposals must commit to code and weight sharing in accordance with CZI's data, publication, dissemination, and other applicable policies and principles. We are not supporting models trained on Protected Health Information, and any human data in the training datasets must be de-identified. Grantees will be subject to applicable

privacy and security requirements and guidelines to ensure that cluster utilization occurs in a safe and compliant manner.

Our Resource

CZI's cluster is a leading-edge resource optimized for the performance and scalability required by AI workloads, thus providing the flexibility and power required to accelerate innovation. The cluster is a CoreWeave GPU cloud-based cluster, designed for advanced workloads, featuring 1,024 Nvidia H100 GPUs in a DGX SuperPod configuration with access to VAST fast data storage optimized for AI and machine learning training at scale. The CoreWeave GPU cloud architecture features a fully managed Kubernetes environment that provides bare metal access for performance and the ability to scale workloads optimally. The cluster uses Vast Data's flash memory (QLC) NVMe SSDs-based storage system to eliminate data access and storage bottlenecks across the AI and data lifecycle. The cluster is actively used to develop and improve cell models on CZI's [prototype virtual cells platform](#).

This RFA aims to support your most audacious biomedical science powered by AI/ML, with the smallest GPU allocation request to be considered at 96 GPUs. This grant is for an allocation of CZI's GPU resource (minimum request of 96 GPUs). **This is an in-kind award; there are no cash funds, financial contributions, or fees of any kind associated with the award.**

ELIGIBILITY

- Applications will be accepted from **domestic** (all collaborating and lead institutions based in the United States) nonprofit organizations, including public and private institutions, such as colleges, universities, hospitals, laboratories, and governmental agencies. For-profit organizations are **not** eligible to apply. All resource-allocation grants will be awarded to institutions, not individuals.
- An organization may submit more than one application (i.e., multiple applicants from the same organization may apply).
- Each application should designate one Principal Investigator (PI) as the Coordinating Principal Investigator (Coordinating PI). Applications may also designate other PIs (Co-PIs; maximum of three). The Coordinating PI will act as the administrative contact between CZI and all other PIs, if any listed on the proposal. The Coordinating PI must submit the application on behalf of all Co-PIs. The Coordinating PI must be affiliated with the institution submitting the application.
- All PIs and Co-PIs must hold a faculty position or equivalent independent investigator status at a **domestic (U.S.-based)** non-profit college, university, medical school, or other non-profit research facility located within the United States. Note that non-tenured and tenured faculty are eligible to apply.
- PIs may serve as Coordinating PI on multiple applications provided that the proposed scope of work for all applications is non-overlapping.
- Projects must propose to utilize at least 96 GPUs on the CZI cluster during the large training phase, with the expectation to scale up over time supported by a clear and well-defined plan.

- Meta employees, including employees of any subsidiary Meta entities and employees of the Chan Zuckerberg Initiative, LLC, are not permitted to apply.
- CZI reserves the sole right to decide if an applicant and applicant organization meet the eligibility requirements.

CZI suggests that you consult your institution to determine eligibility to apply for this grant. For questions about eligibility for this grant or the application process, please contact us before the proposal deadline at sciencegrants@chanzuckerberg.com. Deadline extensions will not be granted.

APPLICATION DETAILS

All applications must be completed and submitted through CZI’s online grants management portal at <https://apply.chanzuckerberg.com>. Applicants are recommended to familiarize themselves with this portal well in advance of the application deadline. Detailed application instructions are available on the [Chan Zuckerberg Initiative website](#) and the [grants management portal](#). **Institutional sign-off by the applicant institution is required as part of the application.**

Key Dates: Applications will be accepted on a rolling basis and will be reviewed at several points during the open application period that runs from January 9, 2025, to June 18, 2025. **All application deadlines and dates are based on Pacific Time.**

Important: Use of the CZI cluster may be fully allocated well before the close of the application period on June 18, 2025. **Once the cluster is fully allocated, the application portal will close, and no further applications will be accepted.** CZI reserves the right to stop accepting applications at any time. Given the rolling nature of this RFA, applicants are encouraged to plan accordingly and submit their applications early.

January 9, 2025	Application portal opens
January 15, 2025 February 19, 2025 March 19, 2025 April 16, 2025 May 21, 2025 June 18, 2025	Submitted applications will be downloaded for consideration at 1 p.m. Pacific Time on the third Wednesday of every month (as listed) until cluster access is allocated. CZI reserves the right to stop accepting applications at any time.
June 18, 2025	Application portal closes at 1 p.m. Pacific Time (PT); CZI reserves the right to stop accepting applications at any time.

Award period and start date: Applications should be submitted for projects with a maximum duration of one year. Resource-allocation award period and start date will be dependent on project scope and will be determined by CZI and the grantee once the project is selected for an award.

Resource Allocation: This grant is for an allocation of CZI's GPU resource (minimum request of 96 GPUs). **This is an in-kind award; there are no cash funds, financial contributions, or fees of any kind associated with the award.**

SELECTION PROCESS

CZI adheres to our [core values](#) in both proposal selection and evaluation of progress. CZI will evaluate all applications for scientific merit through independent expert peer review. Final decisions will be made by CZI staff in consultation with our scientific advisors. There is no expectation of any specific number of awards, and CZI reserves the sole right not to recommend the resourcing of any applications. CZI does not provide feedback on decisions for proposals that are not selected for support.

Selection of grantees will be based on:

- Significance and potential impact of the science proposed, including alignment with CZI's virtual cell research theme and overall mission.
- The expertise and capacity of the proposing team for addressing the proposed project.
- Potential of the grantee to contribute to and benefit from a collaborative interdisciplinary network.

REPORTING AND PROGRESS

Production and availability of intermediate outputs that demonstrate progress is a key mechanism by which CZI evaluates the progress and impact of a project.

- Grantees will share progress, updates, and learnings openly and collaboratively.
- Artifacts from projects, including code produced for software tools, will be made available to the scientific community writ-large via open-source licensing.

POLICIES

- Resources from this award are intended to support research activities in pursuit of building models that accelerate the pace of scientific discovery. Grants are made to organizations to support the work of the Coordinating PI, and reasonable flexibility on how these resources are utilized is allowed, provided that the award is used to support research activities related to the project. **This is an in-kind award; there are no cash funds, financial contributions, or fees of any kind associated with the award.**
- The grantee will be required to submit a final technical report that is due within one month of concluding use of the CZI cluster. Specific deliverable requirements will be outlined in the award notification. As mentioned above, grantees of selected projects will be required to participate in regular meetings, including an annual science meeting (which may be in person or virtual). The grantee may also be required to submit regular technical updates during the project relating to its model development and training and datasets used for training to assist CZI's ongoing support and monitoring of the activities on the cluster, and to ensure CZI's proper maintenance of the cluster and legal compliance.

- **Ethical Conduct:** CZI advocates for the highest standards of ethical conduct in research. In addition to the requirements of their own countries, principal investigators chosen by grantee to complete a grant-supported project shall adopt procedures for the use of animals in research, for the ethical treatment of human subjects and tissue donors, and for obtaining the written informed consent of any human subjects. CZI views the policies of the National Institutes of Health as strong models for such procedures.
- **Data, Publications, and Dissemination Policies:** To accelerate scientific discovery and collaboration, CZI supports a consent, sharing, and publication policy for open and rapid dissemination of research results, including methods, data and reagents, and a policy for software development that maximizes accessibility, reuse, and shared development. Under rare circumstances, exceptions to the above may be considered where there are specific situations that make meeting these goals impossible or counterproductive to the project.
- **Software Code:** CZI requires sharing of software code developed by its grantees generally to be made publicly available. Throughout the project, grantees shall ensure that all code is developed in the open using a code sharing site like GitHub that will be hosted and maintained by CZI, in order to coordinate among other grantees of this RFA, engage the community, and plan cluster allocation. All new code must be released under a permissive open source license (MIT, BSD 2-Clause, BSD 3-Clause, or Apache v2.0), or an alternative community license supporting responsible, open scientific research (as advised by CZI). All releases will be subject to CZI review and approval, for compliance with any applicable legal and regulatory requirements and CZI's responsible AI development guidelines. All pre-existing and derivative code must be licensed under the most permissive license possible, given the licensing terms of the pre-existing code. All analysis packages must be released through the appropriate language-specific package manager (e.g., PyPi for Python, Bioconductor and CRAN for R) with documentation, example data, and interactive demos (e.g., Jupyter notebooks), and the use of Docker or similar container technologies to ensure portability and reproducibility. Software code supported by CZI must be archived for [long-term digital preservation](#) and [citability](#), when applicable.
- **Content and Data Sharing:** CZI is committed to developing and using platforms that disseminate models, data, and benchmarks openly and freely. Any datasets either curated or generated through the proposal must be made as publicly available and easily accessible through an appropriate [data repository](#) as legally permissible, when applicable, under an [Open Definition conformant license](#). Ideally data sets would not include personally identifiable information, but if they do, consent to sharing the data must be obtained. Metadata, documentation, and intended use cases, as appropriate, must be made available under an Open Definition conformant license, preferably CC0 or CC BY/CC BY SA for content that requires explicit attribution.
- **Content Sharing on CZI Platforms:** Models, benchmarks, and datasets developed as part of the RFA will be shared through CZI's public repositories, at CZI's discretion. CZI's repositories, including GitHub, [virtual cells platform](#), [CZ CELLxGENE](#), [CryoET Data Portal](#), or other similar platforms are designed to disseminate resources to the scientific community. All applicants are expected to review and comply with the specific deliverable requirements outlined in the award notification on top of the requirements specified in the Software Code and Data Sharing sections.

- **Publications:** To encourage rapid dissemination of results, any publications related to this awarded work must be submitted to a preprint server (such as bioRxiv, medRxiv, arXiv, or any appropriate preprint repository), at or before the first submission to a journal. Experimental protocols should be made publicly available through a protocol sharing service, such as protocols.io. CZI requires that scientific publications, preprints, and presentations that result from this award acknowledge support from this award.
- **Intellectual Property Rights:** CZI does not require assignment of ownership to any data, published results, or any other intellectual property that results from the work supported by these grants but will have the same rights generally granted to others. CZI supports and promotes policies that enable results and technologies to have the broadest reach and impact. To this end, all newly developed software must be made available through permissive licenses as described more fully above. Other technology and intellectual property rights (such as patents) must be made freely available for all academic and non-commercial use, and where intellectual property rights are commercialized, they must generally be subject to non-exclusive commercial licenses that enable broad availability and dissemination.

CONFIDENTIALITY

All submitted applications that are not selected will be kept confidential, except (1) as necessary for our evaluation or to comply with any applicable laws; and (2) to the extent that the application is made public or available to others without a duty of confidentiality through no fault of CZI or our partner organizations. Notwithstanding the foregoing, certain information, including brief summaries of the proposed projects, project metrics, and types of organizations that have applied for an award, may be made publicly available in aggregate form. Proposals that are selected for an award may be made publicly available and/or shared with other grantees or collaborators. Application materials will not be returned to applicants.

Personal data: Where we ask for personal data of individuals in grant applications, please only submit personal data that you have a right to provide. We will use and store any personal data collected through the application process for grant-related purposes (e.g., administering the grant, analyzing and improving our grant practices). The Chan Zuckerberg Initiative Foundation and Chan Zuckerberg Initiative, LLC (collectively “CZI”) will be the “data controllers” for any such personal information, and the data may be stored on servers outside of your home country, including within the United States. If you have any questions or concerns regarding our privacy practices or the collection or use of personal data, you can contact us at privacy@chanzuckerberg.com.

Detailed Application Instructions

To complete and submit an application:

1. Go to <https://apply.chanzuckerberg.com>.
2. Register and/or log in.
3. Click on the **Programs** link in the upper right corner.
4. Find the **Accelerating and Scaling Biological Sciences with AI RFA** and click **More**.
5. Click the green **Apply** button in the upper right corner.
6. **Enter the title** of your application. The project title is limited to 60 characters, including spaces.
7. Complete the sections described below and **submit by no later than 1 p.m. Pacific Time on June 18, 2025. Please note that applications will be evaluated on a rolling basis, and all resources may be allocated before the final deadline of June 18, 2025.**

The application consists of the following sections (called tasks in the grants portal):

Eligibility, Coordinating PI Details, Organization Details, Project Details, Project Proposal, Biosketches for Coordinating PI and Co-PIs

- **Eligibility:** In an effort to ensure the proposal meets the eligibility requirements to apply, this eligibility questionnaire must be completed first. Once you have completed the questionnaire and confirmed your eligibility, the rest of the application will become available. **We recommend that you complete the eligibility questionnaire well in advance of your anticipated submission date. No extensions will be granted.**
 - Is your project designed to scale to at least 96 GPUs? (Yes/No)
 - Is your project supportive of or directly aligned with CZI's [virtual cell](#) research and/or [overall mission](#) to cure, prevent or manage all diseases by the end of the century? (Yes/No)
 - If yes, describe how (maximum of 100 words).
- **Coordinating PI Details:** Complete all fields in this task; **all fields are required.** The information entered should be for the Coordinating Principal Investigator (Coordinating PI), who will be the person submitting the application on behalf of the team. The Coordinating PI will take responsibility for managing the group collaboration and be the administrative point of contact for CZI and any partners. Information about the Co-Principal Investigator(s) on the proposal should be entered where requested in the Project Details part of the application.
 - Name and email (auto-filled): To edit your name or email, please do so in your account information by clicking your name in the upper right corner and clicking My Account in the dropdown menu.
 - Degree(s)
 - Organization
 - Title/Position
 - Department or equivalent

- Career status: Select early-career (0 to 6 years), mid-career (6+ to 10 years), or neither. **Note: Early- or mid-career status is not required to be eligible for this RFA, although we encourage participation and leadership from early-career researchers.**
- ORCID iD: Enter in format XXXX-XXXX-XXXX-XXXX. ORCID iDs are unique, digital identifiers that distinguish individual scientists and unambiguously connect their contributions to science over time and across changes of name, location, and institutional affiliation. ORCID iDs will be used to streamline reporting in our applications and grant reports to reduce the burden on grantees. For more information and to register, please visit <https://orcid.org/register>.
- **Organization Details:** Complete all fields in this task; **all fields are required.** The information entered should be for the organization of the Coordinating Principal Investigator (Coordinating PI), who will be the person submitting the application on behalf of the team. The Coordinating PI must be affiliated with the organization listed.
 - Organization name/Street address/City/State/Website
 - Type of Organization (Academic, Other Nonprofit, Government, Other)
 - Tax ID: Enter your organization's Employer Identification Number (EIN), as assigned by the Internal Revenue Service in the 9-digit format (XX-XXXXXXX; 10 characters total).
 - Organizational/Administrative Contact: List the name and contact information of the administrative contact to discuss additional information needed, if selected for award.
 - First name, Last name, Title/Position, Email.
 - Signing Official: List the name and contact information of the person authorized to sign on behalf of your organization.
 - First name, Last name, Title/Position, Email.
 - Press Contact / Public Relations Official: List the name and contact information of the person to discuss press releases and media.
 - First name, Last name, Title/Position, Email.
 - Institutional Approval Form: Upload as a single PDF. This [form](#) should be reviewed and signed by a person authorized to sign on behalf of your organization, agreeing to the stated institutional and investigator requirements and commitments on data, resource sharing, and publication policies, as well as endorsing/verifying your application materials and confirming their ability to receive support for the project. The Coordinating PI institution will be responsible for ensuring compliance with all of the terms, including compliance of all partners/subcontract institutions. **These policies are non-negotiable. This form should only be signed if the organization is able to comply with the terms as stated. Note: digital signatures are permitted as long as the document is not encrypted or password-protected.**
- **Project Details**
 - Project Title (auto-filled): Project title is limited to 60 characters, including spaces. If you need to edit your proposal title, navigate to your application summary

- page, click on the three dots to the right of the application title (next to the Preview link) and select Rename from the dropdown menu.
- **Project Purpose:** Summarize your project. Limited to one sentence (maximum of 200 characters, including spaces). **Please use a third-person voice.**
 - **Project Abstract:** Describe your project (maximum of 250 words).
 - **Milestones and Deliverables:** List major milestones and deliverables, and their expected timeline. Be specific and include (where possible) any goals for metrics the project is expected to reach upon completion. Please use a third-person voice (maximum of 250 words).
 - **Co-Principal Investigators:** Select how many Co-Principal Investigators (Co-PIs) will be involved in the project (maximum of three, in addition to the Coordinating PI).
 - If applicable, complete the table with the following information for each Co-PI. **Do not include the Coordinating PI in this section.** You may need to use the scroll bar at the bottom of the table to scroll right to view and complete all fields. Alternatively, you can tab to move through and complete the fields. For each Co-PI, please provide:
 - Co-PI name, Title/Position, Degrees, ORCID iD (format: XXXX-XXXX-XXXX-XXXX), Email, Career status
 - Organization Name and Tax ID (9-digit format: XX-XXXXXXX)
 - **Code Links:** Provide the link(s) to your relevant GitHub repositories, if such repositories currently exist. List up to five maximum using a valid URL in format <https://example.com>.
 - **Frameworks and Software Needs:** Our strong preference is for projects to use the PyTorch framework. If you do not plan to use PyTorch, please specify the Deep Learning or other GPU-bound framework you intend to use (e.g., TensorFlow, XGBoost, MPI v2). Additionally, list any other AI/ML frameworks, software dependencies, or libraries required for your project (maximum of 250 words).
 - **Containerization:** Are your workloads containerized? (Yes/No)
 - If so, what containerization tools or platforms are you using to manage model training and deployment? (maximum of 250 words)
 - **Data Requirements for Model Training:**
 - Describe the datasets you will use for model training, including their size and current storage location (e.g., AWS, on-premises, local devices) (maximum of 250 words).
 - Does your project involve confidential or personal data? (Yes/No) **Note that only de-identified or public data is permitted.**
 - Are there any restrictions or requirements for the public or de-identified data? (Yes/No)
 - **Storage Needs:**
 - How much storage will your project require for input data, models, and other training artifacts? (round up to the nearest TB)
 - Will your model training generate large output files (e.g., trained weights, evaluation logs)? (Yes/No)
 - If so, estimate their size (round up to the nearest TB).

- **Project Proposal:** Upload your project proposal as a single PDF; the font must be 11 point or larger, and margins must be at least one-half inch (top, bottom, left, and right) for all pages (letter size required). Include the following sections:
 - Proposal Body (maximum of 1500 words):
 - Significance and Innovation: Briefly describe the significance of the proposed project, including the model development and training aspects. Highlight specific areas of innovation.
 - Goals/Aims: Detail the major goals of the proposed project, why each is pursued and the anticipated outcomes of each.
 - Approach: Describe the research and development strategy. Describe your approaches to model development, training, and validation. Share any potential pitfalls and your mitigation strategies.
 - Anticipated Outcomes: What will successful completion of the project goals/aims contribute? Both in-field and out-of-field outcomes are of interest. What follow-on research might you anticipate?
 - Project Phases, Timeline, and Resource Requirements: Over a maximum of one year, projects typically progress from smaller experimental stages to larger, more resource-intensive runs. Outline the key phases of your project within this timeframe, including milestones, and specify the number of GPUs or nodes required at each stage.
 - For each phase, outline:
 - a. List the number of GPUs, nodes, and CPU cores you will need.
 - b. Specify the duration of resource usage (e.g., hours, days, months) (maximum of one year).
 - c. Describe how the scale of parameters, data size, or model complexity may evolve over time.
 - d. Address where resource overages might occur and plans for handling constraints if additional capacity is unavailable.
 - Figures/Preliminary Data (optional): Limited to one page, including legends. Figure legends do not count toward the overall word count.
 - References Cited in Your Proposal: No word/page limit; include complete source references.
- **Biosketches for Coordinating PI and Co-PIs:** Upload the biosketches in PDF format for the Coordinating PI and for each of the Co-PIs. Biosketches can be uploaded in a combined single PDF or one PDF for each Co-PI; maximum of five pages per biosketch; [NIH](#) format or similar. Do not include any biosketches for any additional collaborators beyond the Coordinating PI and Co-PIs, as listed.

The formatting and component requirements, including word and page limits indicated above, will be enforced by the review team. Any submitted materials that exceed the word and page limits or do not follow the requirements will not be considered during the application review process.

QUESTIONS?

For administrative and programmatic inquiries pertaining to this RFA, please contact sciencegrants@chanzuckerberg.com. For technical assistance with SMaply, please contact support@smapply.io, or while logged into SMaply, click on the information “i” link in the upper right corner and submit a help request ticket.