ICLR 2026 Workshop Proposal
3rd Machine Learning for Genomics Explorations (MLGenX)

Tagline. Bringing together communities at the intersection of machine learning and genomics to explore new directions in data-driven biology. This year's special track focuses on *From reasoning to experimentation: closing the loop between AI agents and the biological lab.*

Abstract. Despite rapid advances in data-driven biology, our limited understanding of the biological mechanisms underlying diseases continues to hinder therapeutic innovation. While genomics and multi-omics platforms have generated vast datasets, translating these into actionable biological insights remains an open challenge. At the same time, the emergence of foundation models and AI agents capable of reasoning, planning, and hypothesis generation offers a unique opportunity to reimagine how we approach discovery in biology. The 3rd MLGenX workshop aims to bring together the machine learning, genomics, and biology communities to explore this new frontier. This year's theme, "**From Reasoning to Experimentation: Closing the Loop Between AI Agents and the Biological Lab**," focuses on adaptive, interpretable, and experiment-aware AI systems that learn from feedback and drive biological insight. By fostering interdisciplinary collaboration, benchmark sharing, and open discussion, MLGenX 2026 aims to chart the path toward lab-in-the-loop science and accelerate innovation in biology and drug discovery.

Format. In-Person[1].

1 Workshop Motivation and Description

The main objective of this workshop is to bridge the gap between machine learning (ML) and functional genomics, focusing on the new era of AI-driven target discovery and biology. Since its inception, *MLGenX* has aimed to foster a cross-disciplinary community dedicated to exploring how ML can advance biological research and accelerate therapeutic innovation.

Over the past year, there has been remarkable progress in the development of *foundation models* and *AI agents* for biology. Notably, **Biomni**—presented for the first time at MLGenX 2025 by Jure Leskovec (Huang et al., 2025)—demonstrated how reasoning-based systems can plan and execute computational experiments. Several other efforts (e.g., *BioLab* (Jin et al., 2025a), *STELLA* (Jin et al., 2025b), *PerTurboAgent* (Hao et al., 2025), *GeneAgent* (Wang et al., 2025), and *CellVoyager* (Alber et al., 2025)) have since extended this paradigm, marking a transition from using LLMs as passive analytical tools to treating them as *active scientific collaborators*.

These systems exemplify a growing class of AI frameworks that combine reasoning, planning, and calls to foundation models, making biological discovery more transparent, adaptive, and interactive. In parallel, the community has made substantial progress in developing *benchmarks and datasets* for training and evaluating such reasoning-centric, agentic systems.

Looking ahead to 2026, two critical directions are poised to shape the next phase of the field. First, advancing **active learning and lab-in-the-loop frameworks** (Frey et al., 2025) that enable AI Agents to evolve through experimental feedback—closing the gap between *reasoning* and *experimentation*. Second, deepening our understanding of **causality and mechanistic interpretability** in biological foundation models to generate not only accurate predictions but also meaningful, experimentally testable hypotheses.

[1]The talks and panel discussion will be live on SlideLive.

While AI agents represent a forward-looking vision for autonomous discovery, the progress of this field continues to rely on advances in *foundation models* (e.g., *AlphaGenome* (Avsec et al., 2025)) and traditional ML methods that form the computational backbone of these systems. These tools will remain central to generating hypotheses, guiding experiments, and accelerating discovery. Through this workshop, we aim to highlight these emerging paradigms, promote benchmark sharing, and foster interdisciplinary discussions that connect model reasoning, active experimentation, and biological insight.

1.1 Scopes

This year, the workshop will feature four distinct tracks designed to welcome a diverse array of researchers at the intersection of machine learning and biology: the **Main Track**, the **Special Track on Lab-in-the-Loop and Self-Evolving Systems**, the **Tiny Papers Track**, and a new **AI-Generated Track**[2]. By providing dedicated spaces for applied ML in biology, methodological innovation, early-stage ideas, and AI serving as primary author, MLGenX 2026 aims to highlight the next frontier of target discovery—closing the loop between reasoning, experimentation, and adaptation.

1.1.1 Main Track

Machine Learning and Applications. The main track invites contributions addressing foundational and applied challenges at the intersection of ML, biology, and genomics. It emphasizes methodological innovation—spanning foundation models, AI agents, causality, interpretability, and generalization—while also encouraging biologically grounded applications that connect predictive modeling to scientific understanding and therapeutic impact.

Representative methodological topics include:

- **Foundation models and agentic AI:** scalable training, fine-tuning (SFT, RLHF, GRPO, PRM), and in-context learning for biological and multi-omics data; agentic systems for reasoning and tool-use in biomedical research.
- **Causality and mechanistic interpretability:** causal discovery, counterfactual modeling, and biologically grounded representations to reveal mechanism rather than correlation.
- **Generalizability and uncertainty quantification:** developing models that robustly transfer across biological domains, modalities, and experimental settings; quantifying uncertainty in predictive and generative modeling for reliable biological inference.

Applications. We also welcome domain-focused contributions that leverage machine learning for biological discovery and target identification. Example applications include:

- **Design of regulatory sequence elements:** ML for DNA, RNA, and cell or gene therapeutics; sequence-to-function modeling for optimizing UTRs, codons, and regulatory motifs; AI-assisted CRISPR and RNA design.
- **Perturbative biology:** modeling cellular responses to genetic or chemical perturbations; integrating multimodal perturbation readouts (transcriptomic, proteomic, or phenotypic) to understand molecular mechanisms and generalization across contexts.
- **Cellular communication and tissue organization:** causal representation learning to infer cell states, cell–cell interactions, and community structure from multi-omics and spatial data.

[2]All submissions, including those for which AI serves as the primary author, should follow the Policies on Large Language Model Usage at ICLR 2026.

- **Dynamic system simulations:** modeling and simulating multicellular perturbations and disease dynamics; digital twins for treatment optimization and early clinical insight.

1.1.2 Special Track on Lab-in-the-Loop and Self-Evolving Systems

This track focuses on systems that learn through interaction with experimental feedback—moving from predictive models to adaptive, experiment-aware agents capable of autonomous discovery. It centers on lab-in-the-loop frameworks, active learning, and self-evolving agentic systems that directly interact with experimental platforms or simulated environments.

Key topics include:

- **Active learning with experimental feedback:** models that adapt based on wet-lab or simulation outcomes, optimizing experimental design and discovery cycles.
- **Lab-in-the-loop architectures:** integration of AI reasoning systems with robotic labs, automated pipelines, and real-time data streams.
- **Self-evolving agentic frameworks:** autonomous agents capable of iterative hypothesis generation, testing, and refinement using foundation models and multi-agent coordination.

Together, these topics aim to connect the digital and physical layers of biological research, advancing toward autonomous, self-improving scientific systems.

1.1.3 Tiny Papers Track

The Tiny Papers track provides a venue for short, high-impact contributions and early-stage ideas. It is designed to lower the entry barrier for new participants and highlight creative directions, intermediate breakthroughs, or conceptual discussions in a concise two-page format[3].

1.1.4 AI-Generated Track

In addition to the above, MLGenX 2026 introduces a new AI-Generated Track, which welcomes papers substantially authored or co-authored by AI systems *under strict human oversight and ethical review*. Papers in this track should align with either the *Main Track* or the *Special Track* themes (but not the Tiny Papers Track) and follow the standard ICLR format of up to 8 pages excluding references.

This initiative reflects our vision of *AI for Biology* in its broadest sense—spanning both methodological innovation (e.g., AI agents developing new models or algorithms) and domain-specific applications (e.g., perturbation analysis, single-cell modeling, or DNA/RNA design). Its goal is to promote responsible exploration of AI systems not merely as analytical tools, but as active scientific collaborators capable of hypothesis generation, experimentation, and writing.

(i) Permitted AI Systems. Participants may employ advanced AI agents (e.g., *Biomni*, GPT-based models, *BioLab*, or other open systems) for generating text, formulating hypotheses, and/or performing analyses. Each system must be explicitly disclosed, including model name, version, and a brief description of its role in the research process.

[3]AI-generated papers are not allowed for Tiny papers track.

(ii) Human Oversight and Validation. Each submission must designate at least one *human corresponding author* who assumes full responsibility for the content and ethical compliance. Human authors must verify the technical accuracy, originality, and scientific integrity of all AI-generated content; ensure that no private, confidential, or harmful material is included; and carefully proofread, fact-check, and approve the final manuscript before submission.

(iii) Disclosure and Ethics Requirements. To ensure transparency and reproducibility, all AI-generated papers must include:

- **AI Contribution Disclosure:** a short statement detailing the roles of AI and human contributors across the research process (generation, analysis, writing, and visualization).
- **Responsible AI Statement:** a one-paragraph statement (not counted toward the page limit) describing the broader impact, ethical safeguards, and adherence to the ICLR Code of Ethics.

(iv) Review and Evaluation. To further align with the workshop's theme of human–AI collaboration, and to ensure fairness and rigor, all AI-generated submissions will undergo a **dual-review process**:

- **Human reviewers:** at least one domain expert will evaluate novelty, clarity, scientific validity, and the contribution to the ML–biology interface.
- **LLM-based reviewer**[4]**:** an advanced large language model (e.g., GPT-5) will provide an independent structured review assessing coherence, reasoning quality, factual accuracy, and ethical compliance.

An **Area Chair (AC)** will synthesize both review streams and make the final recommendation, taking into account human judgment and AI-assistant feedback. This process will be monitored closely to evaluate the feasibility and trustworthiness of hybrid human–AI peer review in future scientific workflows.

Workshop organizers will conduct additional screening for disclosure completeness and safety compliance. This track aims to explore a new frontier in human–AI collaboration, modeling transparency and accountability while encouraging creative experimentation in AI-assisted scientific authorship.

1.1.5 Additional Highlights

Partnership with *Nature Biotechnology*. Top workshop contributions will be eligible for a fast-track review process at *Nature Biotechnology*. Participation is optional and offered after acceptance notifications (March 1). Selected authors will be invited to expand evaluations and address reviewer comments before transfer to the journal.

BioReasoningChallenge. MLGenX 2026 will host the inaugural *BRChallenge*, designed to benchmark biological AI agents on real-world genomics problems. The topic of this year's challenge will be *Genomics Reasoning* and will include tasks spanning sequence design, perturbation reasoning, and multimodal evidence synthesis to assess the reasoning capabilities of AI agents for target identification. A standardized evaluation protocol with public train/dev splits, hidden validation, and blind test phases will ensure rigor and reproducibility. A public starter kit and benchmark suite will be released prior to the workshop, and top-performing teams will present results and lessons learned during the event.

[4]We follow the Policies on Large Language Model Usage at ICLR 2026.

All invited speakers and panelists have been confirmed.

Time	Speaker	Affiliation	Areas of Expertise
Opening remarks			
9:00 - 9:10	Organizers		
Invited Talks I			
9:10 - 9:40	Charlotte Bunne	EPFL	Digital Twins
9:40 - 10:10	Bo Wang	UHN & U of Toronto	MLLMs for biology
Spotlight I			
10:10 - 10:50	Accepted spotlights		
10:50 - 11:00	Best AI-generated paper		
11:00 - 11:10	Break		
Poster Session I			
11:10 - 12:10	Accepted papers		
12:10 - 13:00	Lunch Break		
Invited Talks II			
13:00 - 13:30	Smita Krishnaswamy	Yale University	Multi-agent systems
13:30 - 14:00	Hani Goodarzi	UCSF & Arc institute	multi-modal FMs
Spotlight II			
14:00 - 14:40	Accepted spotlights		
14:40 - 15:00	BRChallange		
15:00 - 15:10	Break		
Expert Panel Discussion			
15:10 - 16:00	James Zou	Stanford University	AI Agents
	Kyunghyun Cho	NYU & Prescient Design	Lab-in-the-loop
	Kathy Wei	310.ai	Text to design
	Catherine Tong	Isomorphic Labs	GNNs
Round Table			
16:00 - 16:40	Speakers/Panelists/Audience		
Poster Session II			
16:40 - 17:40	Accepted papers		
Closing remarks			
17:40 - 17:50	Organizers		

1.2 Tentative Schedule

The workshop will feature four invited talks from senior and early-career researchers representing diverse areas across machine learning and biology, each lasting 30 minutes (25 minutes presentation, 5 minutes Q&A). To foster interactivity and maximize participation across time zones, we will host **two 60-minute poster sessions**—one in the morning and one in the afternoon—as well as a 50-minute **expert panel discussion** and a 40-minute **interactive round table** with speakers, panelists, and attendees.

Two 40-minute **spotlight sessions** will feature the most notable works accepted to the workshop, with four 10-minute presentations in each. In addition, there will be a dedicated 10-minute segment for presenting the **best AI-generated paper** and a 20-minute segment for the **BRChallenge**, featuring presentations from the two top-performing teams.

This balanced agenda is designed to promote active discussion, encourage cross-disciplinary engagement, and ensure inclusivity for participants joining from different regions.

1.3 Attendees

We expect approximately 450 participants, which amounts to roughly 5% of the total ICLR attendees. This estimation is based on our experience organizing the 1st and 2nd MLGenX workshops at ICLR 2024 and 2025 where 470 and 536 people attended the workshop, respectively [5].

We expect the workshop to attract researchers both from machine learning and biology who are interested in diverse questions ranging from what ML can do for genomics as well as which datasets and questions in the next generation of therapeutics can contribute to advancing fundamental ML research.

1.4 Diversity

Our dedication to diversity, balance, equality, and inclusion is not only represented in the gender balanced organizing team (including senior and junior researchers across several continents and institutions) but also in the invited speakers, panelists, reviewers, and the workshop's organizers:

- The organizing committee includes representation from different affiliations (academia and industry), seniority, geographic location (Europe, USA), gender, and ethnicity.
- The speakers and panelists are chosen from different fields of expertise (both biology and ML), various institutions (both academia and industry), and different seniority levels, genders, and geographic locations.
- One organizer will be involved for the first time in the organization of any workshop.
- The program committee members are also chosen with demographic, expertise, and seniority diversity in mind to minimize the risk of biased judgment of the submissions.
- We offer partial funding for the attendance of two participants from underrepresented countries.
- We offer four distinct tracks—the Main Track, the Special Track on Lab-in-the-Loop and Self-Evolving Systems, the Tiny Papers Track, and the AI-Generated Track—to encourage diverse participation from academia, industry, and early-stage researchers. By welcoming submissions from individuals with a wide range of backgrounds, career stages, and expertise, we aim to create an inclusive environment that fosters collaboration and innovation across all levels of the ML and biology research communities.

1.5 Accessibility, Advertisement, and Website

While the workshop will primarily be held **in person**, we will also utilize a range of digital tools to *engage with our online audience* and enhance interactivity throughout the event. These tools include Gathertown for poster sessions and break discussions, as well as virtual Q&As conducted on Zoom. Moreover, our website will serve as a central platform for disseminating the call for papers, promoting the workshop, and providing early access to the planned agenda and talk titles. This enables attendees to make choices about their attendance based on the content schedule in an accessible manner. We will promote the workshop in advance on our website, via our social media channels, and through collaborations with industry and academic partners to attract a diverse community of researchers interested in machine learning for target identification.

[5] The number of attendees is based on the Whova app.

1.6 Related Workshops

While there are areas of overlap with recent workshops, *MLGenX* is uniquely focused on addressing open problems at the intersection of machine learning and genomics. Unlike workshops centered on drug discovery or structural biology, MLGenX emphasizes high-throughput omics techniques—such as single-cell and perturbation data—that play a pivotal role in bridging genomics and machine learning. These datasets provide the foundation for modeling biological mechanisms, target identification, and reasoning-based AI systems for discovery.

Drug Discovery: ICLR workshop on Machine Learning for Drug Discovery, NeurIPS Workshop on Machine Learning in Structural Biology (MLSB), and NeurIPS Workshop on New Frontiers of AI for Drug Discovery and Development (AI4D3).
In contrast to MLDD, MLSB, and AI4D3, which primarily focus on molecular design, MLGenX centers on the earlier and more mechanistic stage of *target identification*. This focus is crucial not only as the first step in molecule design but also for enabling next-generation therapeutic modalities such as RNA/DNA vaccines, cell therapy, and gene therapy. From a machine learning standpoint, MLGenX stands out by emphasizing reasoning-centric frameworks, foundation models for large-scale genomics data, and the integration of causality, interpretability, and generalizability—areas that complement, rather than overlap with, the focus of existing drug discovery workshops.

Computational Biology: NeurIPS Workshop on Generative AI and Biology (GenBio), ICML Workshop on Computational Biology (CompBio 2023), and NeurIPS Workshop on Learning Meaningful Representations of Life (LMRL).
The proposed MLGenX workshop on target identification stands apart from the previous workshops on computational biology because of its distinct focus and specialized objective. In fact, while past workshops in computational biology have explored the broader spectrum of computational techniques applied to various facets of biology, this workshop puts a strong focus on the three specific areas that connect deep learning to the field of target discovery which is so central in healthcare. From a machine learning perspective, MLGenX shares some similarities with representation learning in LMRL and generative models in GenBio. However, this workshop delves into the foundation models, generalizability, and effective methods for achieving *disentangled* representation learning and interpretability —challenging machine learning topics that we plan to thoroughly discuss during this event.

1.7 Tentative Timeline

- **Call for papers:** 15 December 2025
- **Submission deadline:** 30 January 2026
- **Reviewing period:** 4 February - 21 February 2026
- **Notification:** 1 March 2026

2 Workshop Organizers

Fabian Theis - Director of the Institute of Computational Biology and Professor at TUM Mathematics and Life Sciences, Germany (✉ fabian.theis@helmholtz-munich.de)
Fabian conducts research in the field of computational biology. The main focus of his work is the application of machine learning methods to biological questions, in particular as a means of modeling cell heterogeneities on the basis of single cell analyses and also of integrating "omics" data into systems medicine approaches. Since 2013 he has been a Full Professor of biomathematics at TUM, where he holds the Chair of Mathematical Modeling of Biological Systems, and director of the Institute of Computational Biology at the Helmholtz

Zentrum München.
Previously organized workshops: Fabian organized multiple leading workshops on computational biology such as Workshop on Single Cell Genomics meets Data Science (2022) and Workshop on Computational Single Cell Genomics (2019). He also served as a co-organizer of MLGenX 2024 and 2025.

Aviv Regev - Executive Vice President, Genentech, USA (✉ regev.aviv@gene.com)
Aviv is a computational biologist and systems biologist and Executive Vice President and Head of Genentech Research and Early Development in Genentech/Roche. She is a former core member of the Broad Institute of MIT and Harvard and former professor at the Department of Biology of the Massachusetts Institute of Technology. Regev is a pioneer of single cell genomics and of computational and systems biology of gene regulatory circuits. She co-founded and co-leads the Human Cell Atlas project.
Previously organized workshops: She has also played a role in organizing several conferences and workshops, including AI for Science: Progress and Promises (NeurIPS 2022), Human Cell Atlas General Meetings, and the yearly Single Cell Genomics Conference (SCG). Aviv was also a co-organizer of MLGenX 2024 and 2025.

Arman Hassanzadeh - Senior Software Engineer, Google DeepMind, USA (✉ armanihm@google.com)
Arman is a senior software engineer at Google DeepMind. His work is centered around developing multimodal large language models and developing reasoning agentic systems. Prior to Google (DeepMind), Arman was a PhD student at Texas A&M University working with Nick Duffield and Mingyuan Zhou (UT Austin). His primary research interests are generative models, graph machine learning and Bayesian statistics.
Previously organized workshops: Arman was co-organizer of MLGenX 2024 and 2025.

Mihaela van der Schaar - Professor, University of Cambridge, UK (✉ mv472@cam.ac.uk)
Mihaela is the John Humphrey Plummer Professor of Machine Learning, Artificial Intelligence, and Medicine at the University of Cambridge. She leads the van der Schaar Lab and founded the Cambridge Centre for AI in Medicine (CCAIM). Mihaela is a Fellow of the Royal Society and IEEE Fellow. Her numerous honors include the Johann Anton Merck Award, the Oon Prize on Preventative Medicine, an NSF CAREER Award, multiple IBM Faculty Awards, and the IEEE Darlington Award. She previously served as a Turing Fellow at The Alan Turing Institute (2016–2024) and is currently a Spinoza Guest Professor at Amsterdam University Medical Center.
Previously organized workshops: Mihaela served as lead organizer of the Synthetic Data for Empowering ML Research (NeurIPS 2022) and Synthetic Data Generation with Generative AI (NeurIPS 2024) workshops.

Tommaso Biancalani - Senior Director and Distinguished Scientist, Genentech, USA (✉ biancalt@gene.com)
Tommaso is the head of the BRAID department (Biology Research — AI Development) which is part of the AI for Biology and Translation pillar within the Genentech Computational Science organization. The core mission of the BRAID team is to bridge foundational machine learning research with biology, with emphasis on target discovery. Prior to joining Genentech in 2021, Tommaso was at the Broad Institute of MIT and Harvard where he led a team working on the Human Cell Atlas project. Tommaso trained as a theoretical statistical physicist and completed his post-doctoral training at the Carl Woese Institute of Genomics and MIT.
Previously organized workshops: Tommaso served as a co-organizer for MLGenX 2024 and 2025.

Wei Qiu - PhD student, University of Washington, USA (✉ qiuweipku@gmail.com)
Wei is an incoming Assistant Professor in the Department of Electrical and Computer Engineering at Rice University, with a joint appointment in Computer Science, starting in January 2026. She is currently a final-year Ph.D. student at the Paul G. Allen School of Computer Science and Engineering at the University of Washington. Wei received her B.S. in Data Science and Big Data Technology from Yuanpei College, Peking University. Her research focuses on AI for Biomedicine, developing advanced machine learning methods to study aging and age-related diseases, building foundation models for biomedical data, and designing interpretable and trustworthy AI systems for biology and healthcare.

Ehsan Hajiramezanali - Principal AI Research Scientist, Genentech, USA (✉ hajiramm@gene.com)
Ehsan is a principal AI research scientist at Genentech in the DELTA team within the BRAID department. Before that, he was an AI research scientist at AstraZeneca. Ehsan received his PhD from Texas A&M Uni-

versity, working with Xiaoning Qian and Mingyuan Zhou (UT Austin). His research lies at the intersection of machine learning and Bayesian statistics. He is interested in probabilistic methods, generative models, representation learning, relational learning, and multi-domain learning.
Previously organized workshops: ICLR 2023 Workshop on Machine Learning for Drug Discovery, MARBLE 2023 at ECML-PKDD, MLGenX 2024, and AIDrugX at NeurIPS 2024.

3 Invited Speakers and Panelists

Smita Krishnaswamy. Smita Krishnaswamy is an Associate Professor in the Departments of Computer Science (SEAS) and Genetics (YSM) at Yale University. She is affiliated with the Yale Institute for the Foundations of Data Science, the Wu Tsai Institute, and the Yale Cancer Center, and she participates in the programs in Applied Mathematics, Computational Biology & Bioinformatics, and Interdisciplinary Neuroscience. Her lab develops fundamental deep learning and machine learning methods for representing and learning from high-dimensional biological data. These methods are widely used for data denoising, visualization, generative modeling, dynamics modeling, comparative analysis, and domain transfer in datasets arising from stem cell biology, cancer, immunology, and structural biology, among others. Prior to joining Yale, Smita completed her postdoctoral training in the Department of Systems Biology at Columbia University, where she focused on learning computational models of cellular signaling from single-cell mass cytometry data. She received her Ph.D. in Electrical Engineering and Computer Science from the University of Michigan, where her research focused on algorithms for automated synthesis and probabilistic verification of nanoscale logic circuits. Smita's contributions have been recognized with several honors, including the NSF CAREER Award, a Sloan Research Fellowship, and the Blavatnik Fund for Innovation.

James Zou. James Zou is an Associate Professor of Biomedical Data Science at Stanford University, with courtesy appointments in Computer Science and Electrical Engineering. He is also a member of the Stanford AI Lab and serves as the Faculty Director of the university-wide Stanford Data4Health initiative. His research focuses on developing machine learning methods that are reliable, human-compatible, and statistically rigorous, with broad applications in human health and disease. Algorithms and tools developed by his group are widely used across the technology and biotechnology industries. James received his Ph.D. in Computer Science from Harvard University in 2014. He was previously a member of Microsoft Research, a Gates Scholar at the University of Cambridge, and a Simons Fellow at the University of California, Berkeley, before joining Stanford in 2016. His work has been recognized with numerous honors, including the Sloan Research Fellowship, the NSF CAREER Award, and AI research awards from Google, Amazon, and Adobe.

Kyunghyun Cho. Kyunghyun Cho is a Professor of Computer Science and Data Science at New York University and the Glen de Vries Professor of Health Statistics. He also serves as the Executive Director of Frontier Research at the Prescient Design team within Genentech Research & Early Development (gRED). He is a CIFAR Fellow in the Learning in Machines & Brains program and an Associate Member of the National Academy of Engineering of Korea. Kyunghyun has played an active leadership role in the machine learning community, serving as Program Chair for ICLR 2020, NeurIPS 2022, and ICML 2022. He was one of the three founding Editors-in-Chief of the *Transactions on Machine Learning Research (TMLR)* until 2024. He previously worked as a Research Scientist at Facebook AI Research (2017–2020) and as a Postdoctoral Fellow at the University of Montreal under the supervision of Prof. Yoshua Bengio. He received his M.Sc. (2011) and Ph.D. (2014) degrees from Aalto University, supervised by Prof. Juha Karhunen, Dr. Tapani Raiko, and Dr. Alexander Ilin. His contributions to machine learning and natural language processing have been recognized with numerous honors, including the Samsung Ho-Am Prize in Engineering (2021). He continues to seek a balance among machine learning, natural language processing, and life—though, as he admits, almost always fails to do so.

Charlotte Bunne. Charlotte Bunne is an assistant professor at EPFL in the School of Computer and Communication Sciences (IC) and School of Life Sciences (SV) and a member of the Swiss Institute for Experimental Cancer Research (ISREC). Before, she was a PostDoc at Stanford with Jure Leskovec and

completed a PhD in Computer Science at ETH Zurich working with Andreas Krause and Marco Cuturi. During her graduate studies, she was a visiting researcher at the Broad Institute of MIT and Harvard hosted by Anne Carpenter and Shantanu Singh and worked with Stefanie Jegelka at MIT. Her research aims to advance personalized medicine by utilizing machine learning and large-scale biomedical data. Charlotte Bunne's interdisciplinary research has won several (best paper) awards. Charlotte has been a Fellow of the German National Academic Foundation and is a recipient of the ETH Medal.

Bo Wang. Bo Wang is the Chief Artificial Intelligence Scientist at the University Health Network (UHN) and holds the Canada CIFAR AI Chair at the Vector Institute. His research spans machine learning and computational biology, with a particular focus on developing integrative and interpretable algorithms for biomedical applications. Bo received his Ph.D. in Computer Science from Stanford University, where his work focused on machine learning methods for computational biology, emphasizing integrative cancer analysis and single-cell data modeling. He has extensive industrial research experience at several leading biotechnology and pharmaceutical companies, including Illumina and Genentech. His long-term research goal is to advance machine learning models that improve predictive accuracy and clinical decision support, enabling precision medicine tailored to patients' unique clinical and genomic characteristics.

Hani Goodarzi. Hani Goodarzi is an Associate Professor in the Department of Biochemistry and Biophysics at the University of California, San Francisco (UCSF). He received his Ph.D. in Quantitative and Computational Biology from Princeton University and has led his research group at UCSF since 2016, working at the intersection of machine learning and cancer biology. Hani's lab integrates computational and experimental approaches to investigate how cancer cells reprogram RNA-mediated regulatory networks. In several studies, his team has used machine learning and genomics tools to reveal how RNA-encoded information is altered by cancer cells to drive disease progression, uncovering novel avenues for therapeutic targeting. His contributions have been recognized with numerous honors, including the Vilcek Prize for Creative Promise, the AACR–MPM Transformative Cancer Research Award, the Martin and Rose Wachtel Award in Cancer Research, and designation as an American Cancer Society Scholar. His group brings together students and researchers from diverse backgrounds to tackle fundamental challenges in the life sciences.

Kathy Y. Wei. Kathy Wei, Ph.D., is an ML researcher and co-founder of *310.ai*, a company developing an AI operating system for life sciences. She received her Ph.D. and completed her postdoctoral training in computational protein design at Stanford University with Dr. David Baker. At Amgen, Kathy applied structure prediction, loop and linker modeling, stabilization, and affinity modulation to support therapeutic programs across immunology, oncology, and cardiometabolic disease. She also established and led an interdepartmental effort to launch the AmgenFold platform for large-scale protein structure prediction. As a scientific entrepreneur, Kathy and her team at 310.ai are building the infrastructure to harness the convergence of generative AI and biology—an inflection point they believe will redefine the future of life sciences. She is a multiple patent holder, and her research has been published in leading journals including *Science*, *PNAS*, and *eLife*.

Catherine Tong. Catherine Tong, Ph.D., is a Senior Research Scientist at Isomorphic Labs, where she works at the intersection of machine learning and computational biology to advance drug discovery. She recently completed her Ph.D. at the University of Oxford, where she worked on human activity recognition under the supervision of Dr. Nicholas Lane. During her doctoral studies, she was also a visiting scholar at the Cambridge Machine Learning Systems Lab at the University of Cambridge. Catherine's research combines expertise in machine learning systems, representation learning, and biological data analysis, bridging algorithmic innovation with real-world biomedical applications.

4 Acknowledgment

In addition to our core organizers, we received invaluable support and guidance from several researchers and collaborators. We are deeply grateful to our previous-year organizers, including Mengdi Wang and Sara

Mostafavi, who served as co-organizers of *MLGenX 2025*, and Maria Brbić and Eric Nguyen, who served as co-organizers of *MLGenX 2024*. Their contributions have helped shape the vision, organization, and growth of the workshop series over the years.

We would also like to appreciate Gabriele Scalia and Aïcha BenTaieb for their insightful input, ongoing collaboration, and continued engagement. Gabriele will also lead the *BRChallenge* upon acceptance of the workshop. We extend our sincere thanks to Barbara Cheifet, Chief Editor of *Nature Biotechnology*, for her generous help and guidance in establishing the partnership with the journal.

Finally, we would like to thank the more than **100 program committees** who generously contributed their time and expertise to support the review process in previous years. Their dedication and constructive feedback have been instrumental in ensuring the scientific quality and success of the MLGenX workshop.

5 Program Committee

We have assembled a diverse program committee for the workshop, representing different levels of seniority, fields of expertise, gender, and geographic background to ensure balanced perspectives in the reviewing process. We will ensure that the review load does not exceed *three papers per reviewer*, allowing for constructive and detailed feedback to authors.

If the workshop is accepted, we will announce the call for papers by *15 December 2025*, with a submission deadline of *30 January 2026*. Reviews will be returned to authors by *1 March 2026*. All submissions must represent novel work; previously published papers will not be considered, and reviewers will be instructed to flag and mark such cases accordingly.

Conflict of Interest and Review Policy. The workshop will use the *OpenReview* platform to manage submissions and the reviewing process. All reviews will follow a **double-blind policy**, where both reviewers and authors remain anonymous during the evaluation period. Conflicts of interest will be automatically detected and managed through OpenReview's COI system and verified by the organizers. We will ensure that reviewers and authors are not from the same current institution and have not shared an institutional affiliation within the past five years. In addition, authors and reviewers will be required to declare any personal or professional relationships that could lead to potential bias.

Review Process

All submissions will be managed through the *OpenReview* platform. Accepted papers will be published on the workshop website and designated as *non-archival*, allowing authors to submit extended versions elsewhere in the future. In addition, selected papers—if the authors choose—will be invited to participate in a *fast-track review process* with *Nature Biotechnology*, where they can expand their evaluations and address reviewer feedback prior to transfer.

We will tailor the reviewing process to each track to ensure fairness and rigor:

- **Main and Special Tracks:** Each submission will be reviewed by *three independent reviewers*. An *Area Chair (AC)* will also read the paper and synthesize reviewer feedback to make the final recommendation.
- **Tiny Papers Track:** Each submission will receive *two reviews*. Due to the short format, there will be no AC assigned to this track.
- **AI-Generated Track:** Each submission will receive a *dual review*, consisting of at least one human reviewer and one *LLM-based reviewer*. The AC will evaluate both review streams—human and LLM—to

make the final recommendation. This hybrid review process will be closely monitored to assess the reliability and fairness of human–AI collaboration in peer review.

The program committee members have been selected to cover expertise across machine learning, genomics, computational biology, and AI systems, ensuring broad and balanced evaluation of submissions across all tracks. Below is the tentative list of PC members (reviewers):

- Z. Liu (BeiGene)
- A. Banerjee (Illumina)
- S. Sankarapandian (Calico Labs)
- M. Wang (Purdue University)
- H. Sun (Stanford University)
- N. Gandhi (Dartmouth College)
- O. Tastan (Sabanci University)
- Y. Jiang (Yale University)
- S. Sivanandan (Insitro)
- M. Jindal (Trine University)
- J. Shah (PathAI)
- L. Jiang (EPFL)
- M. Rohbeck (DKFZ)
- L. Cao (Pfizer)
- R. Lyu (CMU)
- A. Wu (MIT)
- R. Qureshi (MD Anderson Cancer Center)
- A. Saadat (EPFL)
- S. Sadhuka (MIT)
- C. Wan (St. Jude)
- M. Koido (University of Tokyo)
- S. Batista (UCLA)
- M. Jia (University of Pittsburgh)
- C. Yuan (CUNY)
- L. Li (Rice University)
- A. BenTaieb (Genentech)
- S. Mourragui (Ensocell)
- S. Niyakan (Texas A&M University)
- H. Huang (Genentech)
- Y. Rosen (Stanford University)
- U. Hazra (Georgia Institute of Technology)
- P. Ghari (Apple)
- A. Xie (MSK Cancer Centre)
- T. Ucar (AstraZeneca)
- S. Jayasundara (Purdue University)
- A. Hashemi (Harvard University)
- A. Agrawal (Parspec)
- H. Jeong (MIT)
- Y. Fu (Calico Labs)
- A. Lin (Atomic AI)
- Y. Liu (Ohio State University)
- X. Huang (Calico Labs)
- W. Guo (UC Davis)
- C. Ye (UC Berkeley)
- R. Littman (Genentech)
- Z. Lu (Genentech)
- B. Zhang (Rice University)
- C. Hu (Mayo Clinic)
- W. Guo (UCLA)
- P. Avdeyev (UT Southwestern Medical Center)
- A. Andersson (Genentech)
- X. Zhang (University of Minnesota)
- L. Lorch (ETH Zürich)
- S. Ghosal (Broad Institute)
- P. Mohseni (Texas A&M University)
- S. Nair (Genentech)
- D. Cakmakci (McGill University)
- E. Nguyen (Stanford University)
- C. Wang (MIT)
- G. Murtaza (Brown University)
- R. Zhang (University of Chicago)
- A. Lobley (University of London)
- Y. Chen (UMD)
- L. Mao (Georgia Institute of Technology)
- N. Wan (Nuanced Health)
- T. Yu (EPFL)
- K. Ni (CMU)
- N. Janakiraman (Google)
- S. Joshi (Columbia University)
- S. Thapa (Intel)
- Z. Pan (Purdue University)
- S. Muller (IBM)
- Y. Wang (ByteDance)
- Y. Yang (University of Michigan)
- A. Feller (UT Austin)
- A. Turcan (CMU)
- R. Rastogi (UC Berkeley)
- X. Liu (Princeton University)
- A. Jain (Salesforce)
- A. Tseng (Amazon)
- Y. Yu (City University)
- G. Mishra (Merck)
- Y. Lee (Genentech)
- Y. Xiao (UCLA)
- S. Maleki (Genentech)
- Y. Wang (UCB Biosciences)
- X. Tu (University of Washington)
- Y. Annadani (TUM)
- N. Diamant (Stanford)
- S. Banerjee (Amazon)

- W. Connell (Transcripta Bio)
- Y. Jin (Harvard University)
- Y. Wu (CUNY)

References

Samuel Alber, Bowen Chen, Eric Sun, Alina Isakova, Aaron J Wilk, and James Zou. Cellvoyager: Ai compbio agent generates new insights by autonomously analyzing biological data. *bioRxiv*, pp. 2025–06, 2025.

Žiga Avsec, Natasha Latysheva, Jun Cheng, Guido Novati, Kyle R Taylor, Tom Ward, Clare Bycroft, Lauren Nicolaisen, Eirini Arvaniti, Joshua Pan, et al. Alphagenome: advancing regulatory variant effect prediction with a unified dna sequence model. *bioRxiv*, pp. 2025–06, 2025.

Nathan C Frey, Isidro Hötzel, Samuel D Stanton, Ryan Kelly, Robert G Alberstein, Emily Makowski, Karolis Martinkus, Daniel Berenberg, Jack Bevers III, Tyler Bryson, et al. Lab-in-the-loop therapeutic antibody design with deep learning. *bioRxiv*, pp. 2025–02, 2025.

Minsheng Hao, Yongju Lee, Hanchen Wang, Gabriele Scalia, and Aviv Regev. Perturboagent: A self-planning agent for boosting sequential perturb-seq experiments. *bioRxiv*, pp. 2025–05, 2025.

Kexin Huang, Serena Zhang, Hanchen Wang, Yuanhao Qu, Yingzhou Lu, Yusuf Roohani, Ryan Li, Lin Qiu, Junze Zhang, Yin Di, et al. Biomni: A general-purpose biomedical ai agent. *bioRxiv*, pp. 2025–05, 2025.

Ruofan Jin, Yucheng Guo, Yuanhao Qu, Ming Yang, Qirong Yang, Linlin Chao, Chun Shang, Yi Zhou, Ruilai Xu, Ziyao Xu, et al. Biolab: End-to-end autonomous life sciences research with multi-agents system integrating biological foundation models. *bioRxiv*, pp. 2025–09, 2025a.

Ruofan Jin, Zaixi Zhang, Mengdi Wang, and Le Cong. Stella: Self-evolving llm agent for biomedical research. *arXiv preprint arXiv:2507.02004*, 2025b.

Zhizheng Wang, Qiao Jin, Chih-Hsuan Wei, Shubo Tian, Po-Ting Lai, Qingqing Zhu, Chi-Ping Day, Christina Ross, Robert Leaman, and Zhiyong Lu. Geneagent: self-verification language agent for gene-set analysis using domain databases. *Nature Methods*, pp. 1–9, 2025.