

Voices

Changing minds: How AI is transforming the life sciences

Artificial intelligence (AI) is changing the way biomedical research is conducted by making it possible to integrate diverse types of biological data and tackle questions that were previously difficult to address. At the same time, its growing use also raises important questions about data quality, reliability, interpretability, and how computational predictions should be combined with biological knowledge and experimental validation. In this Voices piece, researchers from diverse disciplines share their perspectives on how AI is advancing their respective fields and highlight both the transformative opportunities and the key challenges that will shape the future of the life sciences.



Daphne Koller
insitro Inc. and Stanford University, USA

Causal AI won't emerge on its own

The bottleneck in drug discovery has never really been chemistry. It is that we pursue diseases without knowing their cause—the wrong target, in the wrong cell, for the wrong patient. The question that AI is uniquely positioned to answer is a causal one: given an outcome measured at the scale of a whole human, which gene, in which cell type, through which mechanism actually drives it—and would reverse it if perturbed? No person can hold that chain in their head; it spans five or six orders of magnitude, from an ICD code to the molecular state of a single cell. Without AI, it is impossible to integrate evidence across that range and identify genuine causal signal.

Genuine causal leverage in human biology is rare and comes from only two sources—human genetics, which is nature's randomized experiment, and direct perturbation of human cells, which is ours. Each is partial: genetics tells us a gene matters but not how, and perturbation shows how a gene acts but not whether that matters in people. Combined, they let a model reason from outcome to mechanism and select a target with conviction rather than hope.

The most common error is not what AI gets wrong about biology but how we use it. We train models on the observational data we happen to have, then expect causal answers from them. A model trained to predict learns what co-occurs, not what causes what—and no amount of scale converts one into the other. The harder problem is that we lack the causal data to train on. Causality cannot be scraped or inferred; it has to be measured by perturbing a system and recording the result. Generating that data at scale is the real work—and why the goal is still far off. Because the goal was never a virtual cell. It is a causal model of a multi-scale human.



Peilong Lu
School of Life Sciences, Westlake University, China

Synergizing human critical thinking with AI for structural biology

At the heart of structural biology lies the sequence-structure-function paradigm, elucidating the molecular mechanisms driving life. Decades of research have amassed vast repositories of protein sequences and high-quality structures that fuel transformative AI models like AlphaFold for prediction and RFdiffusion and ProteinMPNN for design. These tools transcend the prediction and design of mere protein structures to offer critical insights into how proteins form intricate interaction networks with ions, small molecules, nucleic acids, and polysaccharides. Consequently, AI is uniquely positioned to identify unknown functional macromolecular complexes and their missing components, infer functions via predicted cofactor and ligand interactions, and accelerate the *de novo* design of protein-based therapeutics, diagnostics, and research toolkits. However, while the sequence-to-structure challenge has seen remarkable progress, the structure-to-function gap remains significant. This disparity stems largely from limitations in the quantity, quality, and standardization of functional data.

Current AI systems still struggle to predict dynamic structural ensembles, complex nucleic acid architectures, and ultra-large protein assemblies or to design highly active

enzymes and dynamic molecular machines. Biological function often emerges from context-dependent dynamics and cellular environments that static structure models cannot fully represent. Therefore, human scientific reasoning remains irreplaceable for interpreting complicated biological contexts, recognizing serendipitous discoveries within unexpected results, and designing the critical wet-lab experiments needed to validate uncertain computational predictions. Ultimately, the most vital asset in this AI-driven era is human critical thinking. It empowers researchers to rigorously evaluate which biological questions are suitable for AI, navigate the multitude of computational predictions to select the most promising leads, and remain vigilant against model hallucinations. The future of structural biology depends not on replacing human critical thinking but on synergizing it with AI's power.



Laura Cantini

Institut Pasteur, Université Paris Cité, CNRS, France


Beyond scale: Designing AI for biological discovery

Single-cell sequencing technologies now generate molecular measurements for millions of cells across multiple molecular modalities and spatial coordinates, thereby offering an unprecedented window into cellular organization and functioning. In parallel, AI, particularly foundation models, has become central to extracting patterns from this growing complexity. Yet despite these advances, some of the most fundamental questions in the field remain open: what governs the decision of a cell to commit to a particular fate, and what regulatory logic drives these transitions?

I believe AI is uniquely positioned to address these questions because of the scale and combinatorial complexity of regulatory systems operating across molecular, spatial, and temporal scales. Thousands of biological molecules, interacting nonlinearly across millions of cells and multiple scales, generate a space of biological variation that is difficult to access with classical statistical approaches alone. Modern AI provides the representational capacity needed to learn from these high-dimensional systems. Across the field, AI is increasingly used to integrate heterogeneous data modalities, infer gene-gene interaction networks, and model dynamic cellular processes from large-scale single-cell and spatial datasets. My work focuses on developing methods that explicitly capture these structures to enable the reconstruction of cellular trajectories and the identification of regulatory programs underlying development and disease.

A key emerging challenge is how to fully realize the potential of these models in biological settings. While foundation models already handle batch correction effectively, their ability to learn robust, generalizable representations of cellular dynamics depends strongly on the structure, size, and resolution of the data they are trained on. Single-cell and spatial datasets remain sparse and unevenly sampled, and they often lack multi-modal coverage. As a result, progress toward faithful models of cellular systems will require advances in both experimental design and learning algorithms. A further challenge lies in how biological structure is represented in current models. Many foundation models in biology are adapted from architectures developed for language or vision. While effective for representation learning, these approaches only partially capture the constraints of biological systems, which are defined by gene regulatory networks, spatial organization, and temporal programs. Ignoring this structure risks producing models that are predictive but not reliably grounded in underlying biological mechanisms.

Progress in biomedical AI will therefore depend on moving beyond learning paradigms that assume weak structure in biological data toward models that better reflect the inherent multi-scale organization. By combining modern machine learning with domain-specific principles, we can develop systems that not only predict cellular behavior but also generate mechanistic and testable hypotheses about how biological systems function.



Yuri Pritykin

Department of Computer Science and Lewis-Sigler Institute for Integrative Genomics, Princeton University, USA

Toward predictive immune cell engineering

AI is bringing us closer to a unified model connecting genomic sequence to cellular, multicellular, and organismal phenotypes. The immune system may be where this challenge is hardest and where solving it could matter most. Progress has been fastest in gene regulation, where sequence-to-function problems are relatively well posed and training data are abundant. Yet even there, models remain limited by the lack of data spanning the full diversity of immune cell types and the continuum of cell states across health and disease. Beyond intracellular gene regulation in isolated single cells, the problem is even more challenging. Immune cell fate is shaped not only by the genome, epigenome, and molecular interactions in 3D but also by physical cell-cell contacts, cytokine signaling, and spatial positioning within tissues, all of which change dynamically over time. AI is uniquely positioned to integrate these scales by learning to predict, control, and ultimately engineer immune responses. Reaching that goal requires measurements and perturbations across modalities and spatiotemporal contexts, generated as atlases and shared resources rather than in service of a single hypothesis. Just as deep learning reduced the need for manual feature engineering, foundation models reduce the need for specific hypotheses to justify systematic data generation. But loosening the justification demanded at input requires tightening evaluation at output. Rigorous benchmarking, caveats, publication of negative results, and comparison against simpler models that sometimes perform better are essential; prematurely declaring problems solved only diverts attention from what remains unsolved. Systematic incorporation of immunological domain knowledge is key. Just as formalizing mathematics has enabled recent AI breakthroughs on long-standing open problems, formalizing biological logic will be critical for AI to reason over complex living systems. This rigor is the precondition for using AI to engineer immune cell fate and function and advance diagnostics and therapies for cancer, autoimmunity, infection, and beyond.



Jean Yee Hwa Yang

School of Mathematics and Statistics and Charles Perkins Centre, The University of Sydney, Australia

From descriptive to predictive biology

Advances in genomics, spatial omics, and imaging now enable us to measure biological systems at unprecedented resolution, from molecules and cells to whole tissues. The constraint is no longer data generation; it is integration. We now face the challenge of combining incomplete and heterogeneous observations to understand how tissues are organized and how they change in disease. I believe AI is uniquely positioned to help transform biology from a descriptive science into a predictive one. One exciting opportunity is reconstructing and predicting the molecular state of tissues from partial measurements, such as inferring gene or protein expression from histology and limited spatial data. By learning relationships among tissue structure, cellular organization, and molecular activity, AI can help build virtual tissues that reveal molecular features not directly measured. These models can generate hypotheses, guide experiments, and ultimately help us predict how tissues respond to disease and treatment.

AI will not replace scientists. Its greatest strength lies in learning from vast biological datasets and the scientific literature to identify patterns and relationships. Scientists remain responsible for asking the right questions and interpreting these patterns in context. They must also determine how AI models should be benchmarked and validated, including which datasets, biological settings, and evaluation criteria provide credible evidence that a model is reliable and generalizable. Biological knowledge, scientific creativity, and critical judgment remain essential, particularly when unexpected observations challenge our understanding.

This partnership is changing how I conduct research. I increasingly use AI as a sounding board to test ideas, refine hypotheses, and organize my thinking. It helps me test whether an interpretation is biologically plausible, consider competing explanations, and identify what evidence would distinguish between them. In the long run, AI can be an integral partner in scientific discovery by helping biology move from describing complex systems to predicting their behavior.



Kanaka Rajan
Harvard Medical School, Harvard University, USA

Using AI to reveal how coordination emerges among brains

For me, the central question right now is this: how do brains work together? There is currently a lot of chatter about multi-agent AI and how AI agents can interact with each other to pool resources and solve complex problems. In biological settings, this kind of collaboration is a daily reality. Most of our lives unfold in groups of three to ten people (families, friendship groups, or work teams), and coordinating effectively with them seems to be vital to our mental health. When these close-knit connections dissolve, we experience loneliness, isolation, or the social withdrawal linked to conditions like depression or dementia. The consequences can be profound.

While existing neuroscience techniques can reveal a great deal about a single mind in isolation, what I want to know is the mechanisms behind how brains coordinate in real time within small groups. Standard brain imaging cannot be used to study this scale effectively, as it is limited to scanning one or two individuals at a time in a lab. Conversely, observational studies can track natural group behaviors but reveal nothing about the underlying neural activity. Until now, bridging this gap has been nearly impossible, as we couldn't record inside multiple brains while observing complex group interactions unfold.

AI has changed that. In my research, I use AI "agents": simulated, transparent characters, constrained by the same biology as real brains and trained on real behavioral and neural data from social animals like mice or bats. With several of these agents interacting, we can model group dynamics at scale while tracking the neural dynamics within and between each of the agents simultaneously. This allows us to run experiments that would be impossible or unethical in living animals. We can selectively mute a communication channel, shrink a group, or introduce disruptions to observe precisely how and why coordination breaks down. We can then validate these insights against real brain recordings.

Through this framework, we can map the neural signatures that allow a group to stay connected, adapt to loss, or recover after a conflict. Identifying universal cross-species patterns will help us design better interventions for social isolation. Beyond healthcare, it offers a blueprint for building AI systems that cooperate with humans in ways that are fundamentally safe and trustworthy, rather than merely smart.

In this work, AI is the key to revealing that mechanism. To succeed, these AI agents must solve the problem of coordination using the same biological architecture and limitations as a real brain.



Martin Steinegger
Seoul National University, South Korea

From annotation to mechanism

Understanding the protein universe as a whole has always been a dream, but anyone working with environmental samples such as soil, ocean water, or wastewater encounters the same problem: often, around 80%–90% of identified proteins have no recognizable sequence relatives and remain annotated as hypothetical. In practice, these proteins are frequently excluded from biological interpretation, so we focus on the minority that resembles what we already know and overlook much of what we sequence. AI has changed this. Two proteins can share the same fold despite having almost no detectable sequence similarity. We can now predict accurate structures for many of these unknowns, and comparing shapes rather than sequences reveals relationships for roughly half of them. Viral research benefits disproportionately: viral proteins are notoriously difficult to annotate, yet viruses appear to use a common language of protein folds. This allows us to compare them, place them in an evolutionary context, and find the needle in the haystack. What excites me most, however, is what comes next. Predicting how viral proteins bind host proteins can turn annotation into mechanism: instead of asking whether a sequence comes from a virus, we can ask what it infects and how it interacts with its host. Multimer prediction at scale and interface analysis are central here, and much of this may be possible computationally. Getting there depends on overcoming two related limitations. Complex prediction relies on evolutionary signals shared by both interaction partners, but these signals are often weak between viruses and their hosts. Models also distinguish large structural differences well but struggle with subtle ones, even though host range can turn on only a few

substitutions—the same fold, a handful of altered residues, and a virus may gain the ability to infect humans. The predicted structures can look nearly identical. Resolving differences this small is what separates recognizing a virus from understanding it. Current models capture biology well at the fold level; the next frontier is resolving it at the level of individual residues.



Mihaela van der Schaar

University of Cambridge and the Francis Crick Institute, UK

AI in medicine

For more than a decade, I have worked at the intersection of AI, biology, and medicine, where I have developed approaches including causal machine learning, synthetic data, and digital twins. The biological problem I believe AI is uniquely positioned to tackle in medicine and clinical trials is how to make better decisions earlier, under uncertainty, and for the right patient. Medicine is dynamic: patients deteriorate, recover, relapse, respond differently to treatment, and move through complex care pathways. AI can help us reason about these differences by combining longitudinal and multimodal data, modeling disease trajectories, and asking not only what may happen but what action would change the outcome. This is especially important in clinical trials, where many failures are not failures of the molecule alone but of patient selection, endpoint choice, dose, timing, recruitment, adherence, or assumptions that were never stress-tested. Causal models, synthetic data, and digital twins can help identify which patients are most likely to benefit, which subgroup signals are real, and which trial designs are fragile. But medicine is not just pattern recognition. It is causal, uncertain, biological, ethical, and deeply contextual, and predictive accuracy is not the same as understanding. A trial should learn as evidence accumulates, but continual learning must not become uncontrolled adaptation. AI should support biological, statistical, operational, and regulatory reasoning, not improvise a trial or replace scientific judgment. Every adaptation must be transparent, auditable, prospectively validated, and governed by prespecified rules. This is also why I am excited about my role as chief AI scientist at the Francis Crick Institute. The next frontier is not only AI for medicine but AI that understands biology deeply enough to transform medicine. If AI can help us reason about mechanisms, disease evolution, experimental design, and biological heterogeneity, clinical AI becomes much more than risk prediction: it becomes a bridge from biological discovery to better therapies and better trials.

DECLARATION OF INTERESTS

D.K. is the founder of and an employee of insitro Inc. She and her family members own shares in insitro. M.S. declares an outside interest in Stylus Medicine. M.v.d.S. is a member of the Cambridge Centre for AI in Medicine, which receives funding from AstraZeneca, GSK, Takeda, Boehringer Ingelheim, and McKinsey.