THE ARISING

HEALTH INTELLIGENCE / ARISING HEALTH

FRONTIERS / 06

Can We Simulate a Living Cell?

A guide to biological AI, experimental prediction, and the tests that make a model useful.



Alive. Awake. Aware.

John Griffin, MD / Arising Health

READING EDITION / OCTOBER 25, 2026

How to read this companion

Focused companion based on a 2024 scientific Perspective and a separate clinical-AI feasibility study, checked October 25, 2026. Not a benchmark of current commercial models or evidence that a universal virtual cell exists.

Evidence is attached to the question it answers. Human clinical effects, observational associations, animal experiments, cellular measurements and computational proposals are kept distinct. This reading copy is not individual medical advice.

In this booklet

01 The ambition is larger than a picture

02 What would convince us?

03 Clinical conversation is a different test

04 A practical experiment for The Arising

Follow the evidence

Each chapter identifies the sources supporting its factual claims. Direct primary-source links and review limitations appear in the reference section. Study dates are retained even when older work provides context for a newer story.

01 / The ambition is larger than a picture

Evidence for this chapter: 1. Source details and direct links appear at the end.

A virtual cell is easy to imagine as a beautiful animation. The more important ambition is a model that can represent cellular states and predict how they change after an intervention. Such a tool could help researchers decide which experiments deserve scarce laboratory time.

A 2024 Perspective in Cell describes a vision for an AI virtual cell: a multi-scale, multimodal neural-network-based model linking molecules, cells and tissues across states. Its contribution is a research agenda and design framework. It is not a report that a universally reliable simulator of a living cell has already been built.

That distinction matters whenever the phrase appears in a company announcement. A large dataset, a strong benchmark or an impressive visualization can be valuable without establishing that a model accurately predicts every perturbation in a new biological setting.

The possibility remains significant. If a model can identify useful experiments more efficiently, its value need not depend on reproducing every detail of biology. The right test is whether its predictions improve a defined research task under conditions that resemble actual use.

02 / What would convince us?

Evidence for this chapter: 1. Source details and direct links appear at the end.

The following is an editorial evaluation framework, not a set of clinical findings. First define the task. Predicting a gene-expression pattern, prioritizing a target and forecasting drug toxicity require different evidence. A single accuracy score cannot stand in for all three.

Then ask what was held out. A model evaluated on examples closely related to its training data may do well without learning the

generalization needed for a new cell type, disease state or intervention. A convincing evaluation should make the separation between training and testing understandable.

Next examine the experiment beyond the screen. Did independent laboratory work confirm the prediction? Was the comparator a strong existing method, a simple baseline, or merely random selection? How often did promising predictions fail? Negative results are part of the information needed to judge usefulness.

Finally ask how uncertainty is represented. A model that can identify where it lacks support may be more useful than one that produces a confident output everywhere. That is a proposed design preference; it must itself be evaluated against the task.

03 / Clinical conversation is a different test

Evidence for this chapter: 2, 1. Source details and direct links appear at the end.

The AMIE feasibility study provides a useful contrast. It did not simulate a cell. It evaluated a patient-facing conversational AI system in supervised urgent primary care, with transcripts shared with clinicians. Its evidence concerned feasibility, conversation quality and user experience.

In 98 completed patient encounters, no pre-specified conversation safety stop was required. Supervisors nevertheless recorded one hallucination and added clinical information in five interactions. The system was not tested as an autonomous replacement for care, and zero stops was not zero error.

Why include this comparison in a virtual-cell booklet? Because “AI in medicine” is too broad a label to tell a reader what was demonstrated. A model that supports a biological experiment and a system that helps prepare a clinical visit have different users, failure modes and outcomes.

The shared requirement is to preserve the connection between the claim and the actual evaluation. A computational achievement should not silently become a clinical benefit. A positive user survey should not become a claim about diagnostic safety or patient survival.

04 / A practical experiment for The Arising

Evidence for this chapter: 1, 2. Source details and direct links appear at the end.

For a literature platform, the first useful step is not to claim a digital replica of a patient. It is to make evidence easier to inspect. A proposed workflow could attach each model claim to its dataset, comparator, validation setting and unresolved limitations.

A small evaluation could ask whether AI-assisted literature review retrieves the correct primary sources, preserves denominators, detects conflicting findings and refuses to promote a mechanism into a treatment recommendation. Those are observable tasks with errors that can be counted.

This is an implementation proposal, not a result of the Cell Perspective or a validated Arising Health product. Its value would need to be measured against the existing workflow. The larger vision of virtual biology becomes more credible when the first steps are bounded and testable.

The frontier is not less exciting because a model is incomplete. It becomes scientifically useful when we can specify which parts of the world it represents well, where it fails, and how a better experiment might improve it.

Sources & review scope

1. How to build the virtual cell with artificial intelligence: Priorities and opportunities
December 12, 2024; background framework
Scientific perspective, not a validated universal simulator
PubMed · PMID 39672099
Primary record and complete abstract checked October 25, 2026; full-text/supplement appraisal not completed in this pass.

2. Conversational diagnostic artificial intelligence in ambulatory primary care: a prospective feasibility study
October 24, 2026 journal issue; developer announcement October 8
Prospective single-centre, single-arm feasibility study
PubMed · PMID 42849491
Primary record and complete abstract checked October 25, 2026; full-text/supplement appraisal not completed in this pass.

Illustrations are conceptual editorial art, not photographs of study results or anatomical diagrams. The Horizon and supporting celestial themes are reused from the approved Arising visual collection.

KEEP FOLLOWING THE IDEA

Part of a connected collection

Return to a subject, follow a related booklet, or explore the wider projects. These links organize reading; they do not imply that one scientific finding proves another.

This booklet online →

Booklet directory · all 48 →

Atlas · Single-cell / virtual-cell foundation models →

Arising Health website →

The Arising book →

The Long Game →

Listen & watch →

Related booklets

32 · Molecules We Could Not Previously Make →

33 · A Human Model Before a Human Trial →

45 · Arising Intelligence →

31 · Living Medicines That Respond →

The online review collection is hosted in the shared Atlas. Its existing access settings apply. The public website remains a separate destination.