



# The future of mathematical oncology in the age of AI



Russell C. Rockne<sup>1</sup> ✉, Morten Andersen<sup>2</sup>, Alexander R. A. Anderson<sup>3</sup>, David Basanta<sup>3</sup>, Angela Bentivegna<sup>4</sup>, Sebastien Benzekry<sup>5</sup>, Sergio Branciamore<sup>1</sup>, Sarah C. Brüningk<sup>6,7</sup>, Martina Conte<sup>8</sup>, Farnoush Farahpour<sup>9,10</sup>, Aleksandra Karolak<sup>11</sup>, Alvaro Köhn-Luque<sup>12</sup>, Guillermo Lorenzo<sup>13,14</sup>, Babgen Manookian<sup>1</sup>, Andrei S. Rodin<sup>1</sup>, Lara Schmalenstroer<sup>9</sup>, Juan Soler<sup>15</sup>, Cristian Tomasetti<sup>1</sup> & Konstancja Urbaniak<sup>1</sup>

This perspective article discusses emerging advances at the interface of mechanistic modeling and data-driven machine learning, highlighting opportunities for AI to accelerate discovery, improve predictive modeling, and enhance clinical decision-making. We address critical limitations of current AI approaches and propose a perspective on a future where AI augments mechanistic rigor, clinical relevance, and human creativity under the umbrella of a redefined understanding of Mathematical Oncology.

In May 2025, early-career researchers and established investigators in Mathematical Oncology and artificial intelligence (AI) met to discuss the current and future impact of AI technologies on the field. The focus of the meeting was on immersive learning, spirited debate, and collaborative visioning, moving away from a typical lecture-heavy format. The two-day meeting was held in Siracusa, Sicily, and was structured to help engagement and dialogue, with each day beginning with presentations led by early-career researchers focused on a specific topic, followed by group discussions. Several themes emerged from the discussions that form the basis of this perspective on the future of Mathematical Oncology in the age of AI.

## AI and the future of scientific discovery

A central question was whether AI could eventually outpace human researchers in generating new knowledge and ideas. Several scenarios included whether AI could design viable molecules that don't exist in nature, or discover new physical laws or mathematical conjectures. While some viewed this as speculative, others among us pointed to real-world examples where AI has already begun to do exactly that.

## AI-driven discovery

For instance, in drug discovery, a generative AI platform successfully designed a novel molecule for lung disease, which has since progressed to clinical trials<sup>1</sup>. Recent works have demonstrated how data-driven approaches can elucidate explicit mathematical formulations of the biological mechanisms governing tumor growth and treatment response from longitudinal series of data<sup>2,3</sup>. In continuum mechanics, data-driven computational methods based on regression and neural networks have been demonstrated to enable discovery of mathematical models representing the constitutive mechanical behavior of many different materials<sup>4,5</sup>.

Additionally, an AI model identified new physics governing many-body systems in so called dusty plasma, correcting long-held assumptions. Even in experimental quantum optics, an algorithm called Theseus demonstrated its capacity for explainable AI by designing new, more efficient experimental setups and novel concepts that human physicists have later interpreted and applied<sup>6</sup>.

Meanwhile, in mathematics, AlphaGeometry demonstrated its ability to solve complex geometry problems at a level comparable to a human gold medalist at the International Mathematical Olympiad, showcasing a capacity for complex reasoning and proof generation<sup>7</sup>. Mathematics researchers have integrated AI with proof assistance softwares such as Lean to automate the process of generating and proving new mathematical results<sup>8,9</sup>. Perhaps the most famous example is the generation of protein structure from amino acid sequences generated by AlphaFold<sup>10</sup>. While in most examples, AI is still assisting researchers, rather than working autonomously, it is envisioned that it will soon do so.

At the same time, AI excels at analysis of high-dimensional data, such as the in-depth characterization based on multi-omic profiles: AI and machine learning (ML) are transforming genetic analysis by providing effective tools to handle the large amounts of data generated by genome sequencing. These technologies are proving crucial for detecting and interpreting genetic variants, particularly in oncology. AI algorithms are trained on wide datasets of sequenced genomes to identify the presence of single nucleotide polymorphisms, indels, and structural variations that can discriminate true variants among sequencing errors with a high degree of accuracy, as well as perform analysis and interpretation for clinical use<sup>11,12</sup>.

Many at the conference noted that AI's capabilities are bound by the data it learns from, with severely constrained out-of-distribution reasoning<sup>13</sup>. In domains where data is sparse or non-existent, human

intuition and creativity remain irreplaceable. Think, for instance, of the creative leap of Albert Einstein in formulating the theory of relativity based on the state of knowledge at the beginning of the 20th century. This discussion led to a broader reflection: perhaps the future role of researchers is not to compete with AI, but to use AI as an accelerating tool to work at the frontier of knowledge and science, where data is scarce, uncertainty is high, new paradigms are needed, and where AI will be inherently limited. This acceleration would allow us to check data patterns, make predictions, and test, if not necessarily validate and verify, ideas nearly instantly, and is a way to spend more time understanding the results and asking better questions.

Indeed, the central goal of science is identifying and asking the next question, a task that currently AI is not well equipped to address. The medical sector, and in particular oncology, is one such application field that historically has been strongly supported by human experience and intuition. Changing treatment paradigms limit purely data-driven innovation, as data from contemporary patient cohorts will always be a step behind the next emerging therapy or clinical trial.

### Agentic AI in mathematical oncology

Agentic AI, or the orchestration of multiple large language models (LLMs) as ‘agents’ to execute multi-component workflows, was identified as an advance that necessitates a paradigm shift in how researchers in the field of Mathematical Oncology operate. The group discussed the inevitable development of workflows that mimic the process of mathematical modeling: take experimental or clinical data, construct a family of plausible mechanistic mathematical models, simulate top candidate models, and generate computational codes to reproduce the results in a matter of seconds. Such agentic workflows could be a transformative change to the way Mathematical Oncology research is conducted<sup>14</sup>.

Beyond speed and efficiency, agentic AI offers opportunities for reproducibility and transparency by encoding each step of the modeling pipeline in machine-readable form, potentially reducing human error and facilitating collaborative science across institutions. Moreover, these systems could act as adaptive research assistants, autonomously designing computational experiments, integrating new data streams, or proposing model refinements in response to emerging therapies or clinical outcomes. Early demonstrations in adjacent fields underscore the feasibility of such approaches: multi-agent systems have been used for automated software development, hypothesis generation in the physical sciences, and self-directed optimization of scientific writing<sup>15–17</sup>. In oncology, embedding agentic AI into the modeling pipeline could accelerate the translation of mechanistic insights into clinical strategies, bridging the gap between abstract models and actionable patient-specific predictions.

Nevertheless, challenges remain in ensuring scientific validity, interpretability, and alignment with ethical principles, as agentic AI systems gain autonomy in shaping research trajectories. Addressing these issues will require not only technical safeguards but also new frameworks for oversight, education, and collaborative engagement between human and AI researchers<sup>18</sup>. The importance of integrating AI approaches with mechanistic understanding was also raised as a potential opportunity. By embedding structural and biochemical data into AI systems, researchers can envision multi-scale models that connect atomic-level changes to tissue-level dynamics and patient outcomes. For instance, the concept of temperature is an emergent property that is not defined for an individual molecule. This multi-scale abstraction not only enhances predictive accuracy but also improves interpretability, an essential feature in clinical decision-making. As several participants noted, the challenge with the effective use of agentic AI in Mathematical Oncology is educational: how do we train the next generation to think across scales, disciplines, and leverage AI as an accelerating tool in the development of mathematical models in oncology?

### Mechanistic and data-driven modeling on real-world clinical data

In cancer applications, the very structure of the data often dictates model design. In practice, clinical data may be longitudinal, sparse, irregularly

sampled, and high-dimensional, combining patient covariates, imaging, multi-omics, and circulating biomarkers. While each snapshot can be rich, uneven time points and privacy-driven restrictions on data sharing mean that most longitudinal clinical datasets are “small” in AI terms. Tens, hundreds to at best a few thousand patient trajectories are common, given the data origins from either routinely acquired data, large-scale phase 3 clinical trials, or, in most cases, pooled specific samples across institutions. Merging data from heterogeneous sources, such as different treatment centers, harbors additional challenges regarding data harmonization and comparable protocols, while clinical application of trained frameworks requires extensive evaluation of model robustness and generalization in light of limited sample numbers, and possible data bias. Evolving clinical protocols and instrumentation further add to this challenge, implying that data sparsity and irregular sampling will remain persistent challenges in the clinical data domain without major regulatory and political paradigm shifts.

Purely data-driven models may struggle to digest these complex real-world data realistically, with a possible high chance of overfitting, given the flexibility of these universal function approximators. Mechanistic models, grounded in biological laws, such as tumor growth equations, pharmacokinetics/pharmacodynamics (PK/PD), and tissue mechanics, offer complementary strengths. Physical constraints such as mass conservation and nonnegativity avoid nonsensical extrapolations. Perhaps the most critical feature of mechanistic models is predicting beyond the data, to infer never seen before treatment strategies based on mechanistic knowledge and insights. However, classical mechanistic models often lack the flexibility to integrate unstructured, high-dimensional data or to adapt to unanticipated biology, for example, the emergence of resistant cell populations.

Understanding causal and mechanistic interactions remains a major challenge for AI/ML methods. Temporally dense longitudinal datasets present a significant opportunity to move beyond static associations, enabling the capture of dynamic relationships that evolve over time. However, a particularly critical scenario arises when the assumption of stationarity is violated, a common occurrence in biological systems, where the focus is often on studying transitions between states, such as from health to disease or during stages of development. Identifying the key genetic or epigenetic events that trigger these transitions is especially challenging given their transient and dynamic nature. Importantly, the events that cause state transitions are, by definition, non-stationary and may not align with those exhibiting the greatest differences between steady states. At the molecular level, this is a frequent scenario, with the primary interest often centered on elucidating the sequence of molecular events responsible for transitions such as protein folding/unfolding or activation/inactivation processes.

### Hybrid mechanistic learning

To leverage both mechanistic knowledge and high-dimensional data strengths, hybrid mechanistic learning frameworks overlay data-driven adaptability into an equation-based structure<sup>3,19–24</sup>. Hybrid methods reduce sample complexity by incorporating known dynamics, making irregular and missing data less problematic, and allowing simulation under unseen conditions. Perhaps the most well known example is Physics-Informed Neural Networks (PINNs) that embed governing equations in the loss function to fit sparse data while satisfying known equations<sup>25</sup>.

Similarly, Neural ordinary differential equation (ODE) frameworks learn continuous-time vector fields parameterized by neural networks and can integrate pharmacological insights to discover growth laws from truncated tumor-volume data, improving extrapolation and survival predictions<sup>26</sup>. Operator learning methods, such as DeepONet and Fourier Neural Operators, learn mappings between function spaces, enabling rapid prediction of tumor growth or treatment response in quantitative systems pharmacology<sup>27,28</sup>. In survival analysis, hybrid approaches incorporate mechanistic descriptors, for example, hazard functions derived from tumor growth laws or pharmacokinetic-pharmacodynamic (PK/PD) models, into deep learning (DL) models, improving interpretability and calibration in small-sample contexts<sup>29</sup>.

These hybrid strategies mitigate data sparsity by reducing the sample complexity needed to infer meaningful laws and by making irregular sampling less problematic. For instance, PINNs enforce derivative constraints directly, avoiding the need for evenly spaced grids, and neural ODEs treat each observation as a constraint on the learned vector field. This hybrid thinking forms a conceptual loop: data acquisition informs ML residuals, which refine mechanistic models; these in turn generate verifiable predictions, and the cycle continues. Mechanistic model-driven training, which uses mechanistic models to match patient dynamics as training tools for machine learning models or for cohort generation, circumvents the issue of limited patient training data and potentially allows gaining insights into the patient decisions that impact outcome<sup>30</sup>.

Realizing the potential of these hybrid approaches demands cross-disciplinary fluency. Expertise in biology, mathematics, and machine learning is essential, and training the next generation of researchers to think seamlessly across these domains is as critical as the algorithms themselves. In this integrated view, the most promising advances will likely come from methods that blend the hypothesis-driven rigor of mechanistic modeling, the adaptability of AI, and the geometric insights of topological analysis, all while grounded in the realities of data sparsity.

### Interpretable machine learning methods

There is a tremendous effort to understand how AI models generate specific predictions. While they can yield highly accurate results, the internal decision-making process is often hard to explain, and the models difficult to interpret. The lack of transparency raises concerns about trust, fairness and accountability in critical domains such as medicine and finance, where decisions can directly impact lives. In this context, interpretable and transparent ML techniques, e.g., Bayesian Networks (BNs), offer a distinct advantage<sup>31–33</sup>.

BNs are probabilistic, acyclic graphical models that capture high-order, non-transitive relationships. Unlike neural networks, each relationship is explicitly represented in the topology of the graph. Combined with conditional probability tables that explicitly quantify how variables interact, this structure makes the reasoning process inherently traceable. BNs can represent and test causal relationships, enabling “what-if” scenario analyses that can be especially well-suited for the clinical and scientific settings. This interpretability extends to high dimensional omics data as well; rather than relying on abstract latent features, BNs reveal biologically meaningful and interpretable multiscale nodes and modules.

Broadly applied in oncology, interpretable network modeling extends to the analysis of cellular signaling cascades, therapy-induced state changes, and the emergence of adaptive resistance. Dynamic and temporally resolved BN modeling can reveal the sequence of events that precede cellular transformation, thereby clarifying underlying causal pathways and generating clinically actionable, interpretable hypotheses. Importantly, these data-driven approaches can operate in parallel, or be integrated with mathematical and mechanistic models, facilitating the prioritization of parameters and regimes for targeted investigation in oncological and other biological systems<sup>34,35</sup>.

### Temporal dynamics

Temporal dynamics in biomedical systems and in the mathematical sciences often unfold in high-dimensional spaces where available data are sparse, irregularly sampled, and multimodal. This is particularly evident in cancer applications, where clinical time-series combine imaging, multi-omics, and circulating biomarkers, yet are typically limited to a few hundred or thousand patient trajectories. Privacy-driven restrictions, evolving clinical protocols, and changes in instrumentation ensure that sparsity and irregular sampling remain persistent challenges.

Designing effective modeling frameworks for oncology is uniquely challenging due to the intrinsic nature of the available data. Because tumors evolve according to Darwinian principles, patient responses are dynamic, resulting in clinical time-series that are typically sparse, irregularly sampled, and composed of high-dimensional, multi-modal features such as imaging,

multi-omics, and circulating biomarkers<sup>36,37</sup>. While individual data points can be information-rich, this temporal inconsistency hinders the application of standard deep learning methods that rely on regular sampling. Furthermore, oncology remains a “small data” domain, a reality driven by several persistent factors including data privacy concerns, a lack of dedicated infrastructure for structuring clinical information, significant inter- and intra-tumor heterogeneity, and a constantly shifting landscape of treatments<sup>38,39</sup>. Consequently, even as global data collection efforts expand, the ability of purely data-driven models to accurately capture long-term tumor evolution and treatment effects remains fundamentally constrained by these practical realities.

When confronted with these challenges in the available clinical data, the differences between ML and mechanistic models emerge: ML models are less adept in capturing temporal dynamics. While ML excels at pattern recognition, it often assumes static or dynamic but stationary data – an unrealistic assumption in oncology<sup>40</sup>. Dynamic ML/DL approaches that do not necessarily assume stationarity (such as RNNs, Temporal CNNs, BERT-type Transformers) tend to be sensitive to distributional shifts and computationally demanding for real-world data. Such methods rely on representativeness of prior observations and are therefore inherently limited in extrapolating to new or changing situations often encountered in scientific and clinical research. Mechanistic models, by contrast, are inherently dynamic and hypothesis-driven, grounded in our growing understanding of tumor biology, pharmacokinetics (PK/PD), and systems physiology<sup>41,42</sup>. Digital twins, for instance, aim to integrate clinical data with mechanistic modeling to produce patient-specific, and data-driven predictions<sup>21,43,44</sup>. As such, mechanistic models are ideally suited towards data-sparse regimes and extrapolation, including the modelling of counterfactuals. Notably, many such models (e.g., Latent ODEs / Neural stochastic differential equations (SDEs)) are explicitly designed to handle non-stationarity; oftentimes, it is their *raison d'être*. Importantly, however, mechanistic descriptors inherently remain limited by the assumptions they are grounded in and often lack flexibility to address unstructured data of known mechanistic connection to the trajectory of interest<sup>45</sup>.

While some promising approaches have recently been presented in other application domains, for example physics-informed diffusion models<sup>46</sup>, which embed biological and physical constraints into the generative process, first examples of mechanistic learning solutions, either sequential or as intrinsic combinations have addressed the modelling of tumor response dynamics on imaging as a combination of generative stable diffusion models and mechanistic ODEs describing governing dynamics of tumor size, and the modelling of outcomes, such as progression-free or overall survival<sup>47</sup>. Recent efforts have also focused on combining mechanistic survival models, e.g., based on hazard functions informed by tumor growth laws or PK/PD dynamics, with deep learning models trained on heterogeneous clinical and imaging data to enable biologically plausible extrapolation of survival predictions under different combination treatment scenarios<sup>48</sup>.

A hybrid strategy could reduce dependence on large longitudinal clinical datasets that are costly or infeasible to collect by preserving mechanisms and hypotheses to be studied. Such mechanistic learning solutions maintain an ability to abstract and approximate data regimes of limited mechanistic understanding, harnessing unstructured, deep data.

The same—and additional—challenges apply to all trajectory-indexed data, not just time series: spatial, spatiotemporal, pseudotime, and other biological coordinates. In principle, many time-series-centric statistical/ML methods can be generalized to arbitrary trajectory-indexed settings. However, beyond the issues noted above, especially the computational burden of handling irregular, fine-grained non-stationarity, further obstacles include branching and uncertain trajectories; spatial structure with directionalities, sections, borders, and anisotropy; multiple coordinates/mixed manifolds; spatial periodicity; and assigning orientation without explicit (time) lags. As with temporally indexed data, the field typically attains either fine-grained non-stationarity or scalability—but not both<sup>49–53</sup>.

In our view, the primary thrust should be to develop methodologies that capture interpretable, fine-grained non-stationarity (temporal, spatial, or otherwise trajectory-indexed) without prohibitive computational cost. This can be achieved by implementing dedicated, trajectory-type-specific frameworks that incorporate appropriate biological knowledge and constraints, rather than attempting a general-purpose solution for all trajectory-indexed data out of the box.

### Ethics, uncertainty, and abstraction

Critical ethical and philosophical questions were also discussed. As AI systems increasingly rely on large datasets, for example from national health registries, and begin to be applied to highly temporally resolved patient-specific wearable data, how do we ensure patient autonomy and data sovereignty<sup>54–56</sup>? What frameworks are needed to govern the use of sensitive health data in model training? The protection of sensitive data is an issue even without AI. A solution may be training an AI/LLM on one set of sensitive data, and then apply the trained model elsewhere, for example, via federated learning<sup>57</sup>. This could be a way to benefit from multiple sensitive data sets without directly sharing or combining them.

Further concerns arise as we move toward AI Scientists—AI systems that act not only as tools but as autonomous agents in scientific discovery<sup>58</sup>. Currently, there is a significant implementation gap: while AI systems can generate innovative ideas<sup>59</sup>, they still often lack the ability to follow through with verification, robust experimental design, validation, and interpretation<sup>58,60</sup>. If such systems begin to generate scientific results without clear human oversight, questions emerge regarding scientific accountability and epistemic transparency: who is responsible for the outputs of these AI Scientists?

Finally, as AI systems gain increasing autonomy, there is a growing danger that they could pursue research directions misaligned with ethical standards in healthcare, i.e., autonomy, beneficence, nonmaleficence, and justice<sup>61–63</sup>. Without explicit ethical constraints, autonomous AI agents might prioritize technical novelty or performance metrics over transparency, clinical relevance or patient well-being. Similarly, AI systems might identify therapeutic targets or genetic interventions that are technically feasible but raise concerns about germline editing or reproductive manipulation. In the absence of human judgment and contextual awareness, such systems risk amplifying biomedical research agendas that are not aligned with equity, safety, or patient autonomy. Addressing these risks will require the integration of ethics-by-design principles<sup>64</sup>, domain-specific governance frameworks<sup>65</sup>, and international regulation<sup>66</sup>.

The group also wrestled with the philosophical concept of probability in medicine. Prompted by the provocative article from Spiegelhalter in 2024, the group discussed whether probability exists for single events such as tumor progression<sup>67</sup>. For example, if tumor progression is the predicted outcome with 90% probability, is the prediction wrong if the tumor does not progress? Is the framing of the question as a “right” or “wrong” prediction appropriate in Mathematical Oncology? Is probability a real property of biological systems, or merely a tool for managing uncertainty? While simplification is necessary for modeling, excessive abstraction can obscure critical biological complexity and ignore the reality that probabilities are a necessary means of describing model predictions. The discussion was brought back to practical reality by determining that models must strike a balance between tractability and fidelity, and a recognition that models will never have perfect, complete information, and therefore will inherently have uncertainty, which is best represented as probability.

### Education and training the next generation

A recurring theme was the need to rethink the training of the next generation scientists and in particular Mathematical Oncologists. These include multi-scale thinking and the ability to navigate from molecular biology to systems-level modeling, and from theoretical abstraction to clinical application. Cancer research often requires scientists to move between these different scales and domains, ensuring that the knowledge in one domain can inform learning in others. Realizing these hybrid workflows in practice

demands fluency in biology, mathematics, and machine learning. Training the next generation of mathematical oncologists to think seamlessly across these scales is as critical as the algorithms themselves. Another related, important skill for a successful integration of modeling in cancer research is the competence in interdisciplinary collaboration and effective communication with biologists and clinicians. Scientists must be able to communicate effectively across different fields, each of which has its own terminology, priorities, and methods.

Concerns regarding the misuse of AI tools by students and researchers who lack a foundational understanding were also raised. As AI becomes more accessible, the risk of over-reliance grows, along with the illusion of understanding; the tendency to mistake a fluent, plausible-sounding output for a correct or meaningful one. This can lead to uncritical acceptance of results, especially when users lack the domain knowledge or reasoning skills to evaluate them. Therefore, AI literacy and data science should be another important aspect of training. A reasonable foundation in machine learning, data analysis techniques, and the ethical use of AI tools, more than just familiarity with software, is needed to understand the underlying algorithms and models, as well as their strengths and limitations<sup>68</sup>. AI literacy is not just about using AI tools but critically assessing their outputs. Scientists must be able to evaluate the quality of AI predictions, understand the biases embedded in datasets and AI-generated models, and consider the broader ethical implications of AI decisions. Developing this skill requires courses in both theoretical and practical aspects of AI, paired with hands-on experience in real-world research contexts<sup>69</sup>.

Critical thinking and logical reasoning, in particular, must be deliberately cultivated<sup>70–72</sup>. Just as Mathematical Oncologists are trained to critically evaluate experimental and clinical data by questioning protocol details, choosing appropriate statistical tests, and recognizing selection bias or image manipulation, an analogous set of skills is required to determine when to trust AI-based assessments. Students must learn to examine whether an AI model was trained on appropriate and representative data, question model assumptions, interpret uncertainty, and evaluate the limitations of AI-generated results. This ability also extends to recognizing when AI-generated outputs may reflect biases in the data or in the models themselves. The goal is not just to accept results at face value but to constantly ask, “*What is missing here?*” and “*What assumptions are overlooked?*” Exercises that involve critiquing flawed outputs, whether from a mathematical model or an LLM, can help students build healthy skepticism and scientific rigor. This could involve reviewing AI predictions or evaluating the robustness of models. By engaging with examples of flawed science, whether from human error or model misapplication, students can develop the necessary skills for high-quality research while using AI.

Students must understand that, ultimately, the responsibility for research lies with the human scientist, not the AI tools they use. While AI can aid in processing data and generating insights, it is the researcher who must apply critical judgment, interpret results, and make informed decisions. This awareness of responsibility is key to developing both research maturity and a strong sense of the value they bring to the scientific process. It is essential for students to recognize that the central goal of science is not simply to find answers but to continuously ask the next question. This process of inquiry, which is the drive to explore what comes next, remains at the heart of scientific advancement and, for now, is a uniquely human endeavor. AI may assist in refining or accelerating this journey, but the spark of curiosity, critical thinking, and the formulation of new questions is something only human researchers can contribute<sup>73</sup>.

### Conclusion

In conclusion, Mathematical Oncology stands at a critical inflection point, shaped by the accelerating integration of AI. The convergence of mechanistic modeling with AI-driven methods offers opportunities to overcome long-standing challenges of sparse, irregular, and high-dimensional clinical datasets and simultaneously interpretable, clinically relevant predictions. We advocate for revolutionizing the training of future scientists, emphasizing interdisciplinary skills, critical thinking, and logical reasoning, while

upholding conceptual depth, ethics, and scientific rigor. AI's capacity for accelerating scientific discovery, as evidenced in drug design, the elucidation of biological mechanisms, and the discovery of new physical laws, positions it as a transformative tool. Agentic AI, in particular, offers a paradigm shift by orchestrating multi-component workflows that mimic mathematical modeling processes, thereby enhancing speed, efficiency, reproducibility, and transparency in research.

Furthermore, hybrid mechanistic learning frameworks, which overlay data-driven adaptability onto equation-based structures, are crucial for addressing data sparsity and irregular sampling in clinical data, allowing for more robust predictions and insights. Interpretable machine learning methods, such as Bayesian Networks, provide transparent decision-making processes, which are essential for trust and accountability in medicine. By integrating these areas, we can effectively leverage AI to bridge the gap between abstract models and actionable patient-specific predictions, ultimately accelerating the translation of mechanistic insights into clinical strategies and advancing the field of oncology. The future of Mathematical Oncology lies not in replacing human creativity but in augmenting it, using AI to ask sharper questions, test bold hypotheses, and translate insights into scientific discoveries and ultimately improved patient care.

This perspective was produced with assistance from Microsoft Copilot and ChatGPT 5.0, based on a blog post including a conceptual illustration and summary of a meeting in Syracuse, Sicily in May 2025<sup>74</sup>. The authors provided all prompts, content, and approval of the final version.

### Data availability

No datasets were generated or analysed during the current study.

Received: 30 September 2025; Accepted: 19 January 2026;

Published online: 26 January 2026

### References

- Ren, F. et al. A small-molecule TNIK inhibitor targets fibrosis in preclinical and clinical models. *Nat. Biotechnol.* **43**, 63–75 (2025).
- Podina, L., Ghodsi, A. & Kohandel, M. Learning chemotherapy drug action via Universal Physics-Informed Neural Networks. *Pharm. Res.* **42**, 593–612 (2025).
- Brummer, A. B. et al. Data driven model discovery and interpretation for CAR T-cell killing using sparse identification and latent variables. *Front. Immunol.* **14**, 1115536 (2023).
- Flaschel, M., Kumar, S. & De Lorenzis, L. Unsupervised discovery of interpretable hyperelastic constitutive laws. *Comput. Methods Appl. Mech. Eng.* **381**, 113852 (2021).
- Linka, K. & Kuhl, E. Best-in-class modeling: a novel strategy to discover constitutive models for soft matter systems. *Extrem. Mech. Lett.* **70**, 102181 (2024).
- Krenn, M., Kottmann, J. S., Tischler, N. & Aspuru-Guzik, A. Conceptual understanding through efficient automated design of quantum optical experiments. *Phys. Rev. X* **11** (2021).
- Trinh, T. H., Wu, Y., Le, Q. V., He, H. & Luong, T. Solving olympiad geometry without human demonstrations. *Nature* **625**, 476–482 (2024).
- Yang, K. et al. *Position: Formal Mathematical Reasoning—A New Frontier in AI*, (2025).
- de Moura, L., Kong, S., Avigad, J., van Doorn, F. & von Raumer, J. The lean theorem prover (system description). in *Automated Deduction—CADE-25* 378–388 (Springer International Publishing, 2015). [https://doi.org/10.1007/978-3-319-21401-6\\_26](https://doi.org/10.1007/978-3-319-21401-6_26).
- Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).
- Klemm, F. et al. MSK-ACCESS powered with SOPHiA DDM: Performance analysis of a decentralized MSK-ACCESS solution. *J. Clin. Oncol.* **42**, e15063–e15063 (2024).
- Pozzorini, C. et al. Glinger predicts homologous recombination deficiency and patient response to PARPi treatment from shallow genomic profiles. *Cell Rep. Med.* **4**, 101344 (2023).
- Zhao, C. et al. Is chain-of-Thought reasoning of LLMs a mirage? A data distribution lens. *arXiv [cs.AI]* (2025).
- Swanson, K., Wu, W., Bulaong, N. L., Pak, J. E. & Zou, J. The Virtual Lab of AI agents designs new SARS-CoV-2 nanobodies. *Nat.* 1–8 <https://doi.org/10.1038/s41586-025-09442-9> (2025).
- Yamada, Y. et al. The AI Scientist-v2: Workshop-level automated scientific discovery via agentic tree search. *arXiv [cs.AI]* <https://doi.org/10.48550/ARXIV.2504.08066> (2025).
- McCall, A. & Mccall, A. AI for scientific discovery: Automating hypothesis generation. *Authorea Inc.* <https://doi.org/10.22541/au.175044376.61922933/v1> (2025).
- Alkan, A. K. et al. A survey on hypothesis generation for scientific discovery in the era of Large Language Models. *arXiv [cs.CL]* <https://doi.org/10.48550/ARXIV.2504.05496> (2025).
- Schneider, J. Generative to Agentic AI: survey, conceptualization, and challenges. *arXiv [cs. AI]* <https://doi.org/10.48550/ARXIV.2504.18875> (2025).
- Zhang, R. Z. et al. Personalized predictions of glioblastoma infiltration: mathematical models, physics-informed neural networks and multimodal scans. *Med. Image Anal.* **101**, 103423 (2025).
- Christenson, C. et al. Personalizing neoadjuvant chemotherapy regimens for triple-negative breast cancer using a biology-based digital twin. *NPJ Syst. Biol. Appl.* **11**, 53 (2025).
- Camacho-Gomez, D. et al. Physics-informed machine learning digital twin for reconstructing prostate cancer tumor growth via PSA tests. *NPJ Digit. Med.* **8**, 485 (2025).
- Lorenzo, G. et al. A pilot study on patient-specific computational forecasting of prostate cancer growth during active surveillance using an imaging-informed biomechanistic model. *Cancer Res. Commun.* **4**, 617–633 (2024).
- Lu, J., Deng, K., Zhang, X., Liu, G. & Guan, Y. Neural-ODE for pharmacokinetics modeling and its advantage to alternative machine learning models in predicting new dosing regimens. *iScience* **24**, 102804 (2021).
- Metzcar, J., Jutzeler, C. R., Macklin, P., Köhn-Luque, A. & Brüningk, S. C. A review of mechanistic learning in mathematical oncology. *Front. Immunol.* **15**, 1363144 (2024).
- Zhao, A., Fattahi, D. & Hu, X. Physics-informed neural networks for physiological signal processing and modeling: a narrative review. *Physiol. Meas.* **46**, 07TR02 (2025).
- Laurie, M. & Lu, J. Explainable deep learning for tumor dynamic modeling and overall survival prediction using Neural-ODE. *NPJ Syst. Biol. Appl.* **9**, 58 (2023).
- Hu, Z., Daryakenari, N. A., Shen, Q., Kawaguchi, K. & Karniadakis, G. E. State-space models are accurate and efficient neural operators for dynamical systems. *arXiv [cs.LG]* (2024).
- Mao, R., Wan, L., Zhou, M. & Li, D. Cox-Sage: enhancing Cox proportional hazards model with interpretable graph neural networks for cancer prognosis. *Brief. Bioinform.* **26** (2025).
- Bigarré, C. et al. Mechanistic modeling of metastatic relapse in early breast cancer to investigate the biological impact of prognostic biomarkers. *Comput. Methods Prog. Biomed.* **231**, 107401 (2023).
- Gallagher, K. et al. Mathematical model-driven deep learning enables personalized adaptive therapy. *Cancer Res* **84**, 1929–1941 (2024).
- Hammond, J. & Smith, V. A. Bayesian networks for network inference in biology. *J. R. Soc. Interface* **22**, 20240893 (2025).
- Canonaco, F., Gaudillo, J., Astrologo, N., Stella, F. & Acerbi, E. A guide to Bayesian networks software for structure and parameter learning, with a focus on causal discovery tools. *Front. Syst. Biol.* **5**, 1631901 (2025).
- Gogoshin, G. & Rodin, A. S. Minimum uncertainty as Bayesian network model selection principle. *BMC Bioinforma.* **26**, 100 (2025).
- Zhang, M.-N. et al. Comprehensive review of Bayesian network applications in gastrointestinal cancers. *World J. Clin. Oncol.* **16**, 104299 (2025).

35. Hong, T., Xie, S., Liu, X., Wu, J. & Chen, G. Do machine learning approaches perform better than regression models in mapping studies? A systematic review. *Value Health* **28**, 800–811 (2025).
36. Jiang, P. et al. Big data in basic and translational cancer research. *Nat. Rev. Cancer* **22**, 625–639 (2022).
37. Kazerouni, A. S. et al. Integrating quantitative assays with biologically based mathematical modeling for predictive oncology. *iScience* **23**, 101807 (2020).
38. Bi, W. L. et al. Artificial intelligence in cancer imaging: clinical challenges and applications. *CA Cancer J. Clin.* **69**, 127–157 (2019).
39. Wu, C. et al. A critical assessment of artificial intelligence in magnetic resonance imaging of cancer. *Npj Imaging* **3**, 15 (2025).
40. Lipkova, J. et al. Artificial intelligence for multimodal data integration in oncology. *Cancer Cell* **40**, 1095–1110 (2022).
41. Lorenzo, G. et al. Patient-specific, mechanistic models of tumor growth incorporating artificial intelligence and big data. *Annu. Rev. Biomed. Eng.* **26**, 529–560 (2024).
42. Metzcar, J., Wang, Y., Heiland, R. & Macklin, P. A review of cell-based computational modeling in cancer biology. *JCO Clin. Cancer Inform.* **3**, 1–13 (2019).
43. Hernandez-boussard, T. et al. Digital twins for predictive oncology will be a paradigm shift for precision cancer care. *Nat. Med.* <https://doi.org/10.1038/s41591-021-01558-5> (2021).
44. Abdollahi, H. et al. Theranostic digital twins: Concept, framework and roadmap towards personalized radiopharmaceutical therapies. *Theranostics* **14**, 3404–3422 (2024).
45. Fung, L., Fasel, U. & Juniper, M. Rapid Bayesian identification of sparse nonlinear dynamics from scarce and noisy data. *Proc. Math. Phys. Eng. Sci.* **481** (2025).
46. Bastek, J.-H., Sun, W. & Kochmann, D. M. Physics-informed diffusion models. *arXiv [cs.LG]* (2024).
47. Laslo, D. et al. Mechanistic learning with guided diffusion models to predict spatio-temporal brain tumor growth. *arXiv [cs.CV]* <https://doi.org/10.48550/ARXIV.2509.09610> (2025).
48. Cicolini, J., Barbolosi, D., André, N., Barlesi, F. & Benzekry, S. Mechanistic learning for combinatorial strategies with immunoncology drugs: can model-informed designs help investigators? *J. Clin. Oncol. Precision Med.* 486–491 <https://doi.org/10.1200/PO.19.00381> (2020).
49. Velten, B. et al. Identifying temporal and spatial patterns of variation from multimodal data using MEFISTO. *Nat. Methods* **19**, 179–186 (2022).
50. Zhu, J., Sun, S. & Zhou, X. SPARK-X: non-parametric modeling enables scalable and robust detection of spatial expression patterns for large spatial transcriptomic studies. *Genome Biol.* **22**, 184 (2021).
51. Zhang, K., Feng, W. & Wang, P. Identification of spatially variable genes with graph cuts. *Nat. Commun.* **13**, 5488 (2022).
52. Liu, N. et al. standR: spatial transcriptomic analysis for GeoMx DSP data. *Nucleic Acids Res.* **52**, e2 (2024).
53. Peng, X. et al. Scalable topic modelling decodes spatial tissue architecture for large-scale multiplexed imaging analysis. *Nat. Commun.* **16**, 6619 (2025).
54. Arango-Argoty, G. et al. AI-driven predictive biomarker discovery with contrastive learning to improve clinical trial outcomes. *Cancer Cell* **43**, 875–890.e8 (2025).
55. Swanson, K., Wu, E., Zhang, A., Alizadeh, A. A. & Zou, J. From patterns to patients: Advances in clinical machine learning for cancer diagnosis, prognosis, and treatment. *Cell* **186**, 1772–1791 (2023).
56. Scott, J. M., Stene, G., Edvardsen, E. & Jones, L. W. Performance status in cancer: not broken, but time for an upgrade?. *J. Clin. Oncol.* **38**, 2824–2829 (2020).
57. Pati, S. et al. Federated learning enables big data for rare cancer boundary detection. *Nat. Commun.* **13**, 7346 (2022).
58. Zhu, M. et al. AI Scientists fail without strong implementation capability. *arXiv [cs.AI]* <https://doi.org/10.48550/ARXIV.2506.01372> (2025).
59. Hu, X. et al. Nova: An iterative planning and search approach to enhance novelty and diversity of LLM generated ideas. *arXiv [cs.AI]* <https://doi.org/10.48550/ARXIV.2410.14255> (2024).
60. Padigela, H., Shah, C. & Juyal, D. ML-Dev-Bench: Comparative Analysis of AI Agents on ML development workflows. *arXiv [cs.SE]* <https://doi.org/10.48550/ARXIV.2502.00964> (2025).
61. Mennella, C., Maniscalco, U., De Pietro, G. & Esposito, M. Ethical and regulatory challenges of AI technologies in healthcare: a narrative review. *Heliyon* **10**, e26297 (2024).
62. Farhud, D. D. & Zokaei, S. Ethical issues of artificial intelligence in medicine and healthcare. *Iran. J. Public Health* **50**, i–v (2021).
63. Dankwa-Mullan, I. Health equity and ethical considerations in using artificial intelligence in public health and medicine. *Prev. Chronic Dis.* **21**, E64 (2024).
64. Khan, U. S. & Khan, S. U. R. Ethics by design: A lifecycle framework for trustworthy AI in medical imaging from transparent data governance to clinically validated deployment. *arXiv [cs. CY]* <https://doi.org/10.48550/ARXIV.2507.04249> (2025).
65. Hassan, M., Borycki, E. M. & Kushniruk, A. W. Artificial intelligence governance framework for healthcare. *Healthc. Manag. Forum* **38**, 125–130 (2025).
66. Jobin, A., Ienca, M. & Vayena, E. The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* **1**, 389–399 (2019).
67. Spiegelhalter, D. Why probability probably doesn't exist (but it is useful to act like it does). *Nature* **636**, 560–563 (2024).
68. Robeva, R. S., Jungck, J. R. & Gross, L. J. Changing the nature of quantitative biology education: Data science as a driver. *Bull. Math. Biol.* **82**, 127 (2020).
69. Zha, S. et al. A case study of integrating AI literacy education in a biology class. *Int. J. Artif. Intell. Educ.* <https://doi.org/10.1007/s40593-025-00476-8> (2025).
70. Larson, B. Z., Moser, C., Caza, A., Muehlfeld, K. & Colombo, L. A. Critical thinking in the age of generative AI. *Acad. Manag. Learn. Educ.* **23**, 373–378 (2024).
71. Naqvi, W. M., Ganjoo, R., Rowe, M., Pashine, A. A. & Mishra, G. V. Critical thinking in the age of generative AI: implications for health sciences education. *Front. Artif. Intell.* **8**, 1571527 (2025).
72. Papanephytous, C. & Nicolaou, S. A. Promoting critical thinking in biological sciences in the era of artificial intelligence: The role of higher education. *Trends High. Educ.* **4**, 24 (2025).
73. Kannan, M., Bridgewater, G., Zhang, M. & Blinov, M. L. Leveraging public AI tools to explore systems biology resources in mathematical modeling. *NPJ Syst. Biol. Appl.* **11**, 15 (2025).
74. Rockne, R. et al. *The Future of Mathematical Oncology in the Age of AI* <https://mathematical-oncology.org/blog/future-of-math-oncology.html> (2025).

### Acknowledgements

The authors acknowledge the generosity of the city of Ortigia, Syracuse, Sicily, for providing the meeting venue and refreshments and the Beckman Research Institute at City of Hope for financial support.

### Author contributions

R.C.R. attended the meeting in Syracuse and contributed to and approved the final manuscript. M.A. attended the meeting in Syracuse and contributed to and approved the final manuscript. A.R.A.A. attended the meeting in Syracuse and contributed to and approved the final manuscript. D.B. attended the meeting in Syracuse and contributed to and approved the final manuscript. A.B. attended the meeting in Syracuse and contributed to and approved the final manuscript. S.Be. attended the meeting in Syracuse and contributed to and approved the final manuscript. S.Br. attended the meeting in Syracuse and contributed to and approved the final manuscript. S.C.Bru. attended the meeting in Syracuse and contributed to and approved the final manuscript. M.C. attended the meeting in Syracuse and contributed to and approved the final manuscript. F.F. attended the meeting in Syracuse

and contributed to and approved the final manuscript. A.K. attended the meeting in Syracuse and contributed to and approved the final manuscript. A.K.-L. attended the meeting in Syracuse and contributed to and approved the final manuscript. G.L. attended the meeting in Syracuse and contributed to and approved the final manuscript. B.M. attended the meeting in Syracuse and contributed to and approved the final manuscript. A.S.R. contributed to and approved the final manuscript. L.S. attended the meeting in Syracuse and contributed to and approved the final manuscript. J.S. attended the meeting in Syracuse and contributed to and approved the final manuscript. C.T. attended the meeting in Syracuse and contributed to and approved the final manuscript. K.U. attended the meeting in Syracuse and contributed to and approved the final manuscript. All authors read and approved the manuscript.

### Competing interests

The authors declare no competing interests.

### Additional information

**Correspondence** and requests for materials should be addressed to Russell C. Rockne.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026

<sup>1</sup>Department of Computational and Quantitative Medicine, Beckman Research Institute, City of Hope, CA, USA. <sup>2</sup>Centre for Mathematical Modeling - Human Health and Disease, Roskilde University, Roskilde, Denmark. <sup>3</sup>Department of Integrated Mathematical Oncology, Moffitt Cancer Centre, Tampa, FL, USA. <sup>4</sup>School of Medicine and Surgery, University Milano-Bicocca, Italy; Fondazione IRCCS San Gerardo dei Tintori, Monza, Italy. <sup>5</sup>Centre Inria d'Université Côte d'Azur and Cancer Research Center of Marseille, Institut Paoli-Calmettes, Inserm, CNRS, Aix Marseille University, Marseille, France. <sup>6</sup>Department of Radiation Oncology, Inselspital, Bern University Hospital and University of Bern, Bern, Switzerland. <sup>7</sup>Department of Digital Medicine, University of Bern, Bern, Switzerland. <sup>8</sup>Politecnico of Torino, Turin, Italy. <sup>9</sup>Group of Bioinformatics and Computational Biophysics, University of Duisburg-Essen, Essen, Germany. <sup>10</sup>Institute for Cell Biology (Cancer Research), University Hospital Essen, Essen, Germany. <sup>11</sup>Department of Machine Learning, Moffitt Cancer Centre, Florida, USA. <sup>12</sup>Oslo Centre for Biostatistics and Epidemiology, Faculty of Medicine, University of Oslo, Norway and Department of Medical Genetics, Oslo University Hospital, Oslo, Norway. <sup>13</sup>Group of Numerical Methods in Engineering, Department of Mathematics and CITEEC, University of A Coruña, A Coruña, Spain. <sup>14</sup>Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin TX, USA. <sup>15</sup>Department of Applied Mathematics and Research Unit Modeling Nature (MNat). University of Granada, Granada, Spain. ✉ e-mail: [rockne@coh.org](mailto:rockne@coh.org)