Mechanistic Learning as a combination of Machine Learning and Modelling in Mathematical Oncology

Sarah Brüningk (Department of Radiation Oncology, University of Bern, Inselspital, Switzerland),
Renee Brady-Nicholls (Moffitt Cancer Center, USA),
Kathleen Wilkie (Department of Mathematics, Toronto Metropolitan University, CA)

5.1.2025–10.1.2025

1 Overview of the Field

Mechanistic learning is an emerging paradigm at the intersection of mathematical modelling and machine learning that aims to synergize domain knowledge with data-driven discovery [1, 2, 8]. Within the field of mathematical oncology, this hybrid approach is gaining traction as a powerful framework for building models that are not only accurate and data-efficient but also interpretable, generalizable, and clinically actionable. Cancer research poses unique modelling challenges due to its reliance on limited, heterogeneous data and its demand for translational relevance in high-stakes clinical settings [3]. Mechanistic learning offers a way to bridge the conceptual and technical divide between these established modelling traditions to help address the challenges in cancer research.

In its classical form, mathematical modelling in oncology involves the formulation of equations and/or predefined rules grounded in biological, chemical, and physical principles to simulate complex processes such as tumor growth, metastatic spread, and treatment response [4]. These knowledge-driven models allow researchers to extrapolate and explore future or unobserved scenarios, and to perform *in silico* experiments, ultimately striving to generate mechanistic hypotheses for experimental or clinical validation. The transparency and interpretability of such mechanistic models make them especially valuable for clinical translation [5]. However, a key limitation of these frameworks remains the unavoidable simplifications which have a risk of oversimplification, particularly given sparse or noisy data that may not sufficiently cover all aspects of assumptions made.

In contrast, data-driven models, particularly those based on modern machine learning and deep learning techniques, are designed to uncover complex relationships within large, high-dimensional datasets. Typical examples include imaging data, -omics data, or electronic health record data. These models often function as powerful statistical tools for prediction and classification, excelling in tasks such as tumor subtyping, response prediction, and image segmentation [6]. Nevertheless, they are frequently criticized for being "black boxes" [7] that lack interpretability and biological grounding and their performance depends heavily on the availability of large, high-quality datasets, which are not always feasible to obtain in biomedical research.

Mechanistic learning offers a framework for combining mechanistic and data-driven paradigms in a mutually reinforcing manner. Rather than treating knowledge- and data-driven modelling as isolated or even competing strategies, mechanistic learning conceptualizes a spectrum of integration strategies that leverage the strengths of each. As articulated in Metzcar et al. (2024)[8], this spectrum can be understood as a landscape comprising four main types of combinations: sequential, parallel, extrinsic, and intrinsic.

Sequential approaches involve the stepwise application of one method to support the other. For example, machine learning models may be used to estimate parameters or initial conditions for a mechanistic model, or mechanistic outputs may serve as features for a downstream statistical model. These methods are relatively easy to implement and interpret, though they may face challenges in uncertainty propagation. Parallel approaches treat data-driven and mechanistic models as complementary alternatives applied to the same task. Their outputs can be compared or ensembled, providing a measure of robustness or uncertainty quantification. This is particularly attractive in clinical applications, where confidence in predictions is critical.

Extrinsic combinations utilize one approach to post-process or interpret the output of the other. For instance, machine learning can be used to correct residual errors from a mechanistic model, or mechanistic theory can be employed to validate and explain patterns discovered by data-driven methods. Intrinsic combinations represent the most tightly coupled approaches, embedding domain knowledge directly into the structure or training objective of the learning model. Physics-informed neural networks (PINNs) are a prime example, incorporating differential equations as soft constraints in the model's loss function. Similarly, biologically informed neural networks may encode known regulatory interactions into the model's architecture, thereby constraining its capacity in a structural way.

All of these strategies are not mutually exclusive and are best understood as points along a continuum. Mechanistic learning is thus less a set of fixed tools than a conceptual landscape—one that allows researchers to navigate between extremes of purely theoretical and purely empirical modelling. This view promotes methodological flexibility and innovation, encouraging the design of hybrid models tailored to the nature of the data, the specific research question, and the intended application. Several practical examples underscore the potential of this approach. Neural networks trained as surrogate models can approximate the solutions of computationally intensive PDE systems with orders-of-magnitude faster inference times [9]. Symbolic regression has enabled the automated discovery of governing equations from empirical data [10], while data-

driven residual modelling has improved treatment outcome predictions beyond mechanistic models alone [11]. Hierarchical Bayesian models and digital twins further illustrate how multilevel, multimodal integration is not only possible but already underway in precision oncology [12, 13].

By framing mechanistic learning as a fluid and adaptable interface between data and theory, this perspective fosters interdisciplinary collaboration. It offers a shared modelling language that bridges applied mathematics, systems biology, clinical oncology, and computer science. In doing so, mechanistic learning aligns with the broader goals of computational precision medicine: to move beyond mere statistical association toward a more mechanistic, interpretable, and ultimately actionable understanding of disease processes.

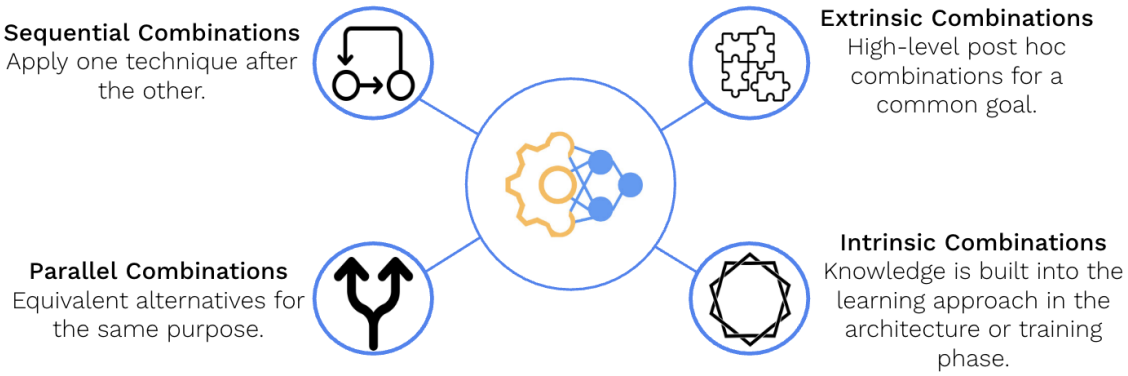

Figure 1: The mechanistic learning landscape highlighting examples of how mechanistic mathematical modelling and data-driven machine and deep learning approaches can be combined in various forms with different levels of computational and conceptual complexity.

The BIRS workshop "Mechanistic Learning as a Combination of Machine Learning and Modelling in Mathematical Oncology" positioned itself at the forefront of this interdisciplinary effort, emphasizing not just technical innovation but also the importance of collaboration across domains. The field stands poised to contribute significantly to the development of robust, interpretable, and personalized computational tools that can enhance clinical decision-making in oncology with the key objective of establishing long-lasting new collaborations across disciplines and geographical location. Given the recent rise in interest in the topic of mechanistic learning, this workshop was designed as a think-tank-type event offering insight into the current landscape of approaches, and to motivate participants to look beyond the current state-of-the-art and be creative in the design of new paradigms or applications in the context of mechanistic learning.

2 Recent Developments and Open Problems - Workshop Expectations

The rapid growth of mechanistic learning as a field is not only driven by technical advances but also by the increasing recognition of its relevance to real-world biomedical challenges. This momentum was clearly reflected in the BIRS workshop, where a uniquely diverse group of researchers came together to explore, shape, and extend the frontier of this emerging domain. Participants came from a wide range of backgrounds—including mathematical modelling, machine learning, systems biology, and clinical oncology—and represented a healthy distribution of career stages, from doctoral students to established faculty. This diversity was essential to capturing the full complexity of mechanistic learning and ensuring that discussions remained grounded in both mathematical rigor and clinical relevance.

At the start of the workshop, participants were invited to articulate their expectations and aspirations (see Figure 2).

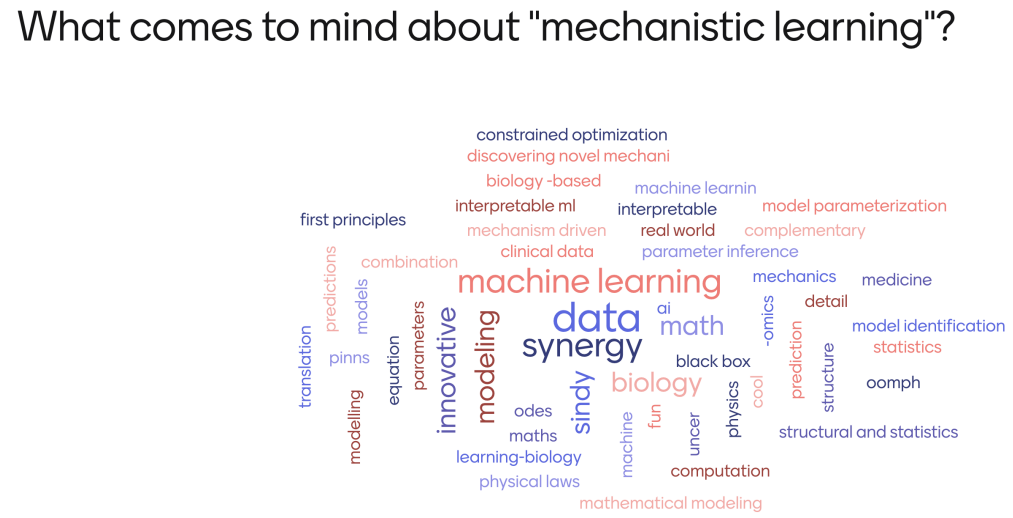

Figure 2: Word cloud obtained in the opening survey regarding the expectations for the upcoming workshop amongst participants.

Many expressed a desire to move beyond siloed expertise and engage in genuinely interdisciplinary dialogue. Some hoped to gain practical exposure to machine learning frameworks, while others aimed to learn how mechanistic modelling principles could inform their data-driven work. Several early-career researchers voiced a clear need for more interpretable, transferable, and clinically grounded models—objectives that mechanistic learning is uniquely positioned to address. There was a shared recognition that neither mechanistic modelling nor machine learning, on their own, can fully resolve the challenges of precision oncology. Instead, participants expressed a collective motivation to explore how combining both approaches might help tackle problems that are currently inaccessible to either field in isolation. This was particularly relevant given the differences in research backgrounds of our participants (see Figure 3).

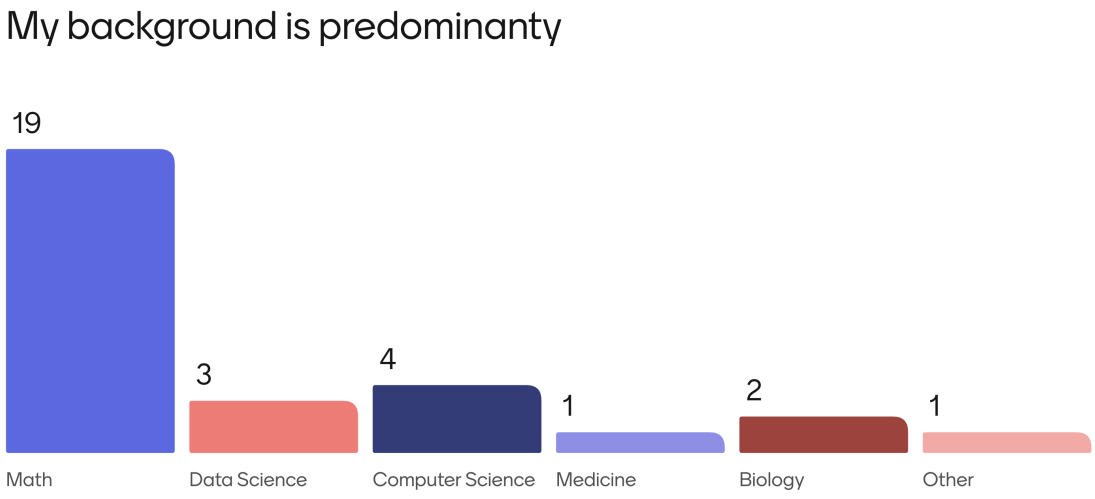

Figure 3: Survey of the background areas represented by our participants.

This spirit of collaboration catalyzed several promising developments. Groups began exploring the integration of domain knowledge into neural architectures, developing strategies for hybrid inference pipelines, and investigating the feasibility of digital twin frameworks for individualized treatment planning. These

projects often took the form of proof-of-concept studies, reflecting the workshop's emphasis on exploration and foundational insight over immediate implementation. Importantly, the discussions also revealed several persistent open problems that the field must address to achieve broader clinical impact.

One pressing challenge is the issue of data scarcity and heterogeneity. While data-driven models excel with large, standardized datasets, real-world biomedical data are often limited, noisy, and inconsistently annotated. Mechanistic learning offers ways to mitigate this, but many participants noted that new strategies are needed to quantify and propagate uncertainty, especially in hybrid models that span both statistical and mechanistic spaces.

A second key area of concern is model interpretability and explainability. Although intrinsic combinations such as physics-informed neural networks are a step in this direction, participants pointed out that more work is needed to ensure that hybrid models are not only accurate but also transparent and understandable to clinicians and regulators. This is particularly important in high-stakes applications such as treatment selection and response prediction, where black-box models may not be acceptable.

Third, the field still faces integration challenges at the methodological level. There is no unified pipeline or framework for building, training, and validating mechanistic learning models. As a result, groups often resort to ad hoc combinations of tools, which may hinder reproducibility and scalability. Developing standardized benchmarks, shared datasets, and open-source tools was highlighted as an urgent priority.

Finally, participants emphasized the importance of sustained interdisciplinary collaboration. Mechanistic learning thrives at the intersection of communities, but continued progress will require ongoing efforts to bridge language, culture, and training differences across fields. Mentorship, inclusive research environments, and follow-up initiatives were all identified as critical to maintaining the momentum sparked during the workshop.

As such, our BIRS workshop not only showcased recent developments in mechanistic learning but also surfaced a set of open challenges that now serve as a collective research agenda. By building on the shared insights and motivations of its diverse participants, the field is poised to make meaningful strides toward interpretable, data-efficient, and clinically relevant modelling in oncology and beyond.

3 Presentation Highlights

This workshop was designed with a predominantly interactive approach, and a selection of presentations included to initially offer some background on the topic, and later to include specific research examples. We kicked off with an opening presentation (by all organizers) to specify the objectives and deliverables for the week. Importantly, participants were given the freedom to identify individually what a "measurable outcome" implied for their group and were offered a variety of options in the form of a blog post, publication, joint grant application, or seminar invitations as possible examples for measurable outcomes. We concluded the introductions with two ice-breaker game-style interactions in which course participants got to mingle and learn about each other's backgrounds, expertise, and interests. The organizers then opened the scientific program with three introductory presentations on the topics of mechanistic mathematical modelling (presenter: R. Brady-Nicholls), data science (presenter: S. Brüningk), and an avenue of meachanistic learning as a combination of the two preceding approaches (presenter: S. Brüningk). Importantly, these presentations laid the groundwork for participants from a multidisciplinary background to obtain brief, yet focused insight regarding the tremendous parallels, but also implementation specific differences for mechanism and data-driven approaches. Given the high information density, participants were also provided with the slides to be able to go through the material in their own time and throughout the workshop as an inspiration and source of useful references.

The remaining time of the workshop was allocated to small group (4-8 participants) discussion and project formulation. Each team was pre-allocated based on different criteria (career stage, background) and predefined preferences for the topics of the participants collected in a preparation survey. In order to ensure a goal-oriented approach and productivity, teams initially designed their outcome goal on day one and pitched these alongside their timeline for the week on Monday. Teams, then continuously provided updates to their progress on a daily basis, followed by a final presentation on Friday.

In addition to interactive group work, we invited five speakers from various career levels, scientific backgrounds, and research focus areas to present ongoing projects in the domain of mechanistic learning for

mathematical oncology applications. These included speakers onsite, as well as a virtual presentation and these presentations were also streamed for all virtual participants as well:

- Jana Gevertz (Professor of Mathematics, The College of New Jersey, USA, onsite): **"Model-driven experimental design with applications to cancer immunotherapy"**. Dr. Gevertz's presentation focused on optimizing experimental designs in cancer immunotherapy through mathematical modelling. She emphasized the importance of identifying model-informative data to ensure unique parameter estimation, thereby enhancing the predictive power of models. Her recent work includes developing frameworks that integrate identifiability analysis to guide data collection strategies, ultimately aiming to improve the design of virtual clinical trials and therapeutic interventions.

- Thomas Yankeelov (Professor, Director of Center for Computational Oncology, University of Texas at Austin, remote presentation): **"Imaging based digital twins for patients who are not average"**. Dr. Yankeelov discussed the development of digital twins—computational models that replicate individual patient tumor dynamics—by integrating advanced imaging techniques with multi-scale mathematical models. His approach aims to personalize cancer treatment by forecasting tumor responses to therapies, thereby optimizing treatment plans for patients whose characteristics deviate from population averages.

- Jana Lipkova (Assistant Professor for digital pathology, University of California Irvine, onsite): **"Bayesian Digital Twins for Radiotherapy Planning in Neuro-Oncology"**. Dr. Lipkova presented a Bayesian framework for creating digital twins to enhance radiotherapy planning in neuro-oncology. By integrating multimodal imaging data, such as MRI and PET scans, with mathematical tumor models, her approach enables the estimation of tumor cell densities and the quantification of uncertainties in tumor infiltration. This methodology facilitates the design of personalized radiotherapy plans that aim to maximize efficacy while minimizing damage to healthy tissue.

- John Nardini (Assistant Professor in the Department of Mathematics and Statistics at The College of New Jersey, onsite): **"Forecasting and predicting stochastic agent-based models of cell migration with biologically-informed neural networks"**. Dr. Nardini explored the use of biologically-informed neural networks (BINNs) to predict the behavior of stochastic agent-based models (ABMs) that simulate cell migration. By training BINNs to learn interpretable differential equations from ABM data, he demonstrated improved forecasting capabilities and the potential to predict system behavior under untested conditions. This approach offers a promising avenue for efficiently exploring complex biological systems where traditional modelling techniques may fall short.

- Lena Podina (PhD candidate, University of Waterloo, onsite): **"Universal Physics-Informed Neural Networks and their Applications"**. Ms. Podina introduced Universal Physics-Informed Neural Networks (UPINNs) as a tool for discovering symbolic differential operators in scenarios with sparse experimental data. By incorporating prior knowledge about the underlying dynamics into the learning process, UPINNs can identify governing equations that describe complex systems accurately. Her work highlights the potential of UPINNs in enhancing model interpretability and reducing computational costs in data-limited settings.

Moreover, throughout the workshop we showcased specific research contributions from individual participants in poster sessions during the coffee breaks. This gave everyone, and in particular more junior researchers, an opportunity to present their work in the context of the ongoing workshop. Given the limitations in space, posters rotated throughout the week, offering everyone a chance to participate. This format was perceived as very engaging and led to fruitful discussions. We are also happy to report a good involvement of our virtual participants, who (where this was desired) were able to actively contribute to the group projects in addition to joining the plenary updates and research presentations. Finally, our mentoring event and matching of early-career and experienced researchers were very well received. In total, ten mentor-mentee pairs were matched across geographical locations and given the opportunity to discuss future career goals and project specifics in our dedicated mentoring session on the first evening to set the scene for further discussion in the upcoming days.

4 Scientific Progress Made

The workshop was designed to foster interdisciplinary collaboration in the emerging field of mechanistic learning. Participants were pre-assigned to diverse groups blending expertise in modelling, machine learning (ML), and clinical science. On the first day, each group selected a mechanistic learning topic, pitched their research idea with anticipated outcomes, and provided daily progress updates. This iterative and interactive structure enabled rapid feedback, adaptation, and cross-pollination of ideas throughout the event. In the following, we summarize the individual projects addressed.

4.1 ML for Model Parameterization

Team: Sara Hamis, David Hormuth, John Metzcar, Maria Monzon Ronda, Bernadette Stolz
This project explored the integration of machine learning (ML) into the parameterization of mechanistic models for glioblastoma (GBM). The team focused on three main topics (see Figure 4): i) using ML to assign tumor-specific parameters, ii) replacing hard-to-model terms (like radiation death rates), and iii) identifying recurrence biomarkers from clinical data. Using topological data analysis on multiparametric MRI data from preclinical glioma models, the group extracted features representing spatial heterogeneity, then applied regression models to predict key parameters like diffusion coefficients. Another effort used ML to replace model terms in radiation response models, leveraging in vitro confluency data. The third topic involved a data-driven approach to predict tumor recurrence based on longitudinal white blood cell counts in GBM patients. The team faced challenges with noisy, sparse data and model identifiability but laid a clear path for future subgroup analyses, model refinement, and potential submission to the annual meeting of the Society of Mathematical Biology.

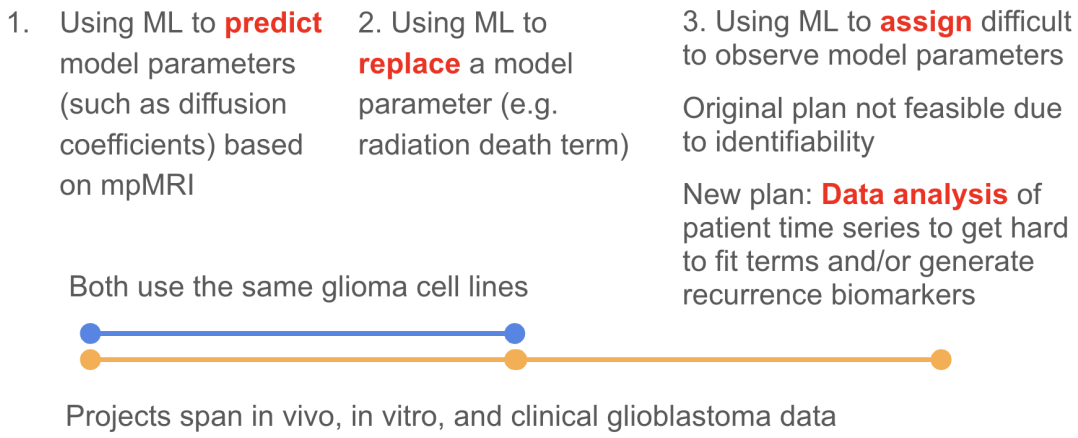

Figure 4: Overview of the addressed topics in the context of using machine learning for model parameterization.

4.2 Mechanistic Filters

Team: Morten Andersen, Russell Rockne, Divyanshu Tak, Jana Lipkova, Matt Faria, Sarah Brüningk, Guillermo Lorenzo, Philipp Altrock
This group developed the Deep-learning Integrated Structural SINDy Mechanistic Filter (DISS-MF) for predicting GBM recurrence (see Figure 5). By combining sparse identification of nonlinear dynamics (SINDy) with deep learning and structural priors, the team aimed to enhance predictive accuracy for spatial recurrence patterns in GBM. Leveraging the Brain Imaging Adaptive Core (BrainIAC) database with over 50,000 brain MRIs, they created a hybrid model capable of learning tumor morphology and drug distribution from imaging. Tasks were distributed among coding, generating concentration maps, and extending models to radiotherapy. The group aimed to produce a research publication and an abstract for SNO, with planned follow-ups and

cross-continental collaboration. Their work reflects a rigorous mechanistic learning approach, fusing data-driven methods with biological interpretability.

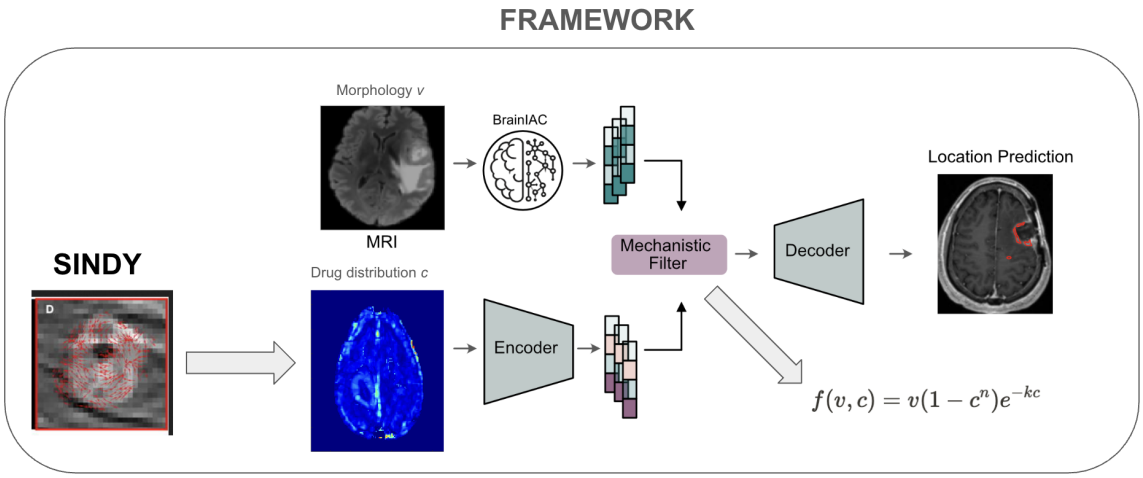

Figure 5: Overview of the mechanistic filter application in the context of a prediction of the location of GBM recurrence using a combination of structural information from MRIs, and a distribution of systemic therapies modelled using a SINDy framework to describe flow fields based on dynamic contrast-enhanced MR imaging.

4.3 Mechanism-Informed Neural Networks (MiNN)

Team: Suzan Farhang-Sardroodi, Xinyang Liu, Jonas Latz, Lena Podina, Adam Stinchcombe, Gilbert Koch
This project targeted pediatric diffuse midline gliomas (DMGs), a highly fatal and poorly understood brain tumor. The team proposed a Mechanism-Informed Neural Network (MiNN) approach, incorporating prior mechanistic knowledge into the network architecture as outlined in Figure 6. Using data from 40 pediatric patients, including MRI scans pre- and post-radiation therapy (RT), the model sought to infer underlying tumor growth dynamics by learning an intervention term that adapts to individual responses. Their mathematical model combined diffusion, logistic proliferation, and a learned radiation intervention term. Challenges included small datasets and maintaining team engagement, but preliminary results showed promising trends. The team plans to expand to 3D modelling, apply real data, and submit a paper by summer 2025, dividing tasks among modelling, coding, and data validation.

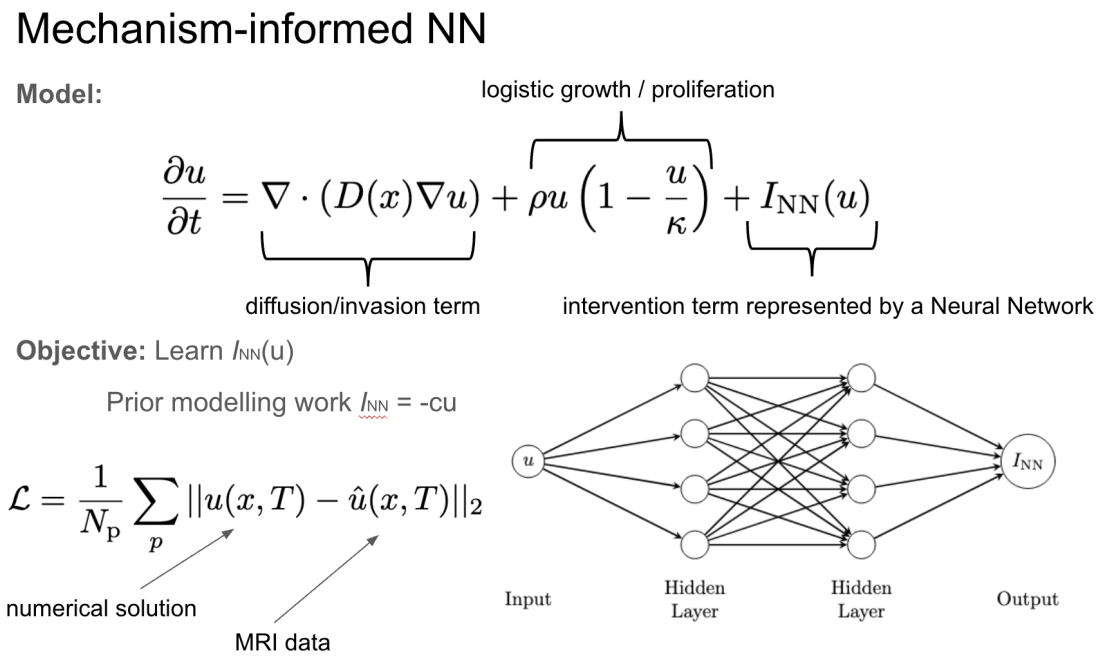

Figure 6: Mechanism informed neural network design comprising a mechanistic 3D tumor growth model driven by cell invasion, and growth, complemented by an approximated network term for treatment effects.

4.4 INSITE – *In Silico* Trials using Simulation and AI

Team: Rebecca Bekker, Renee Brady-Nicholls, Tyler Cassidy, Lisette de Pillis, Jana Gevertz, Harsh Jain, Ali Nabavizadeh

The INSITE project designed a comprehensive Virtual Clinical Trial (VCT) pipeline integrating mechanistic modelling and AI to evaluate cancer therapies more efficiently than traditional trials. Figure 7 shows an overview of the required steps in this pipeline, which includes drug target identification, model development, calibration, validation, virtual patient generation, and trial analysis. At each stage, the team assessed where ML can enhance efficiency or insight—e.g., drug repurposing, symbolic regression for model structure, and clustering for treatment response stratification. Major challenges included data sparsity, model integration, and translating preclinical findings to clinical relevance. The team is preparing a perspective article for Mathematical Biosciences and plans to continue work through blog posts and collaborative publications, emphasizing mechanistic learning's potential to revolutionize oncology trials.

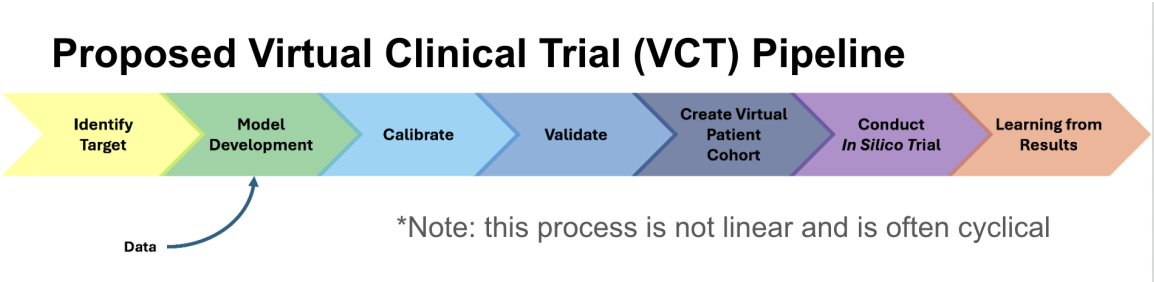

Figure 7: Schematic illustration of the proposed virtual clinical trial pipeline comprising a chain of building blocks from the domains of mechanistic modelling and data-driven AI.

4.5 UDEs vs. UPINNs

Team: Kathleen Wilkie, Alvaro Köhn-Luque, John Nardini, Changin Oh, Gibin Powathil, Erica Rutter, Jeffrey West

This group explored the comparative performance of Universal Differential Equations (UDEs) and Universal Physics-Informed Neural Networks (UPINNs) in modelling tumor dynamics. They specifically contrasted these often interchangeably used definitions (see Figure 8). Using synthetic and biological data, they analyzed the ability of these architectures to learn accurate dynamics under varying assumptions: unknown, partially known, or incorrectly specified model components. Their motivating case was tumor clone competition interactions, with radiation treatment effects. The study highlighted the strengths of UPINNs in enforcing biological structure while allowing flexibility, contrasting with the more unconstrained UDEs. The team will produce a MathOnco blog post and a research article comparing these methods on real and synthetic datasets. Challenges included code standardization (eventually converging on Julia) and collaborative integration across modelling frameworks.

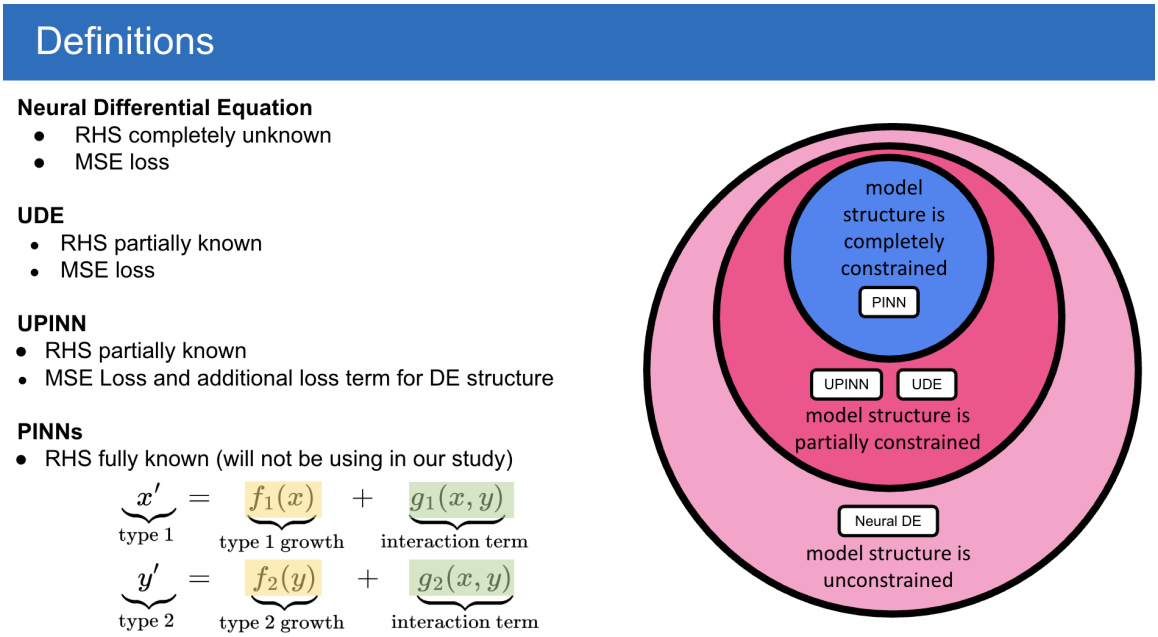

Figure 8: Key definitions provided by the UDE vs. UPINNs group to contrast these key concepts and illustrate their relationships.

4.6 GEnAI and Mechanistic Modelling (GEMM)

Team: Ari Barnett, Daria Laslo, Zhifan Jiang, Heiko Enderling

The GEMM group proposed a synergistic framework that fuses Generative AI (GenAI) with mechanistic tumor growth models in the context of radiation oncology as outlined in Figure 9. The team highlighted how GenAI can aid in tumor segmentation, anatomical prediction, and imputing missing timepoints—enhancing spatiotemporal understanding of tumor evolution. Conversely, mechanistic models provide structure and constraints to guide GenAI outputs. A key concept introduced was probability convolution applied to diffusion models to predict tumor growth uncertainty over time. The team aims to formalize this hybrid approach with proof-of-concept studies and a perspectives article. Challenges remain in defining the quantitative oncologist's role in this ecosystem, but the project reflects a forward-thinking vision for integrating data-driven innovation into personalized oncology workflows.

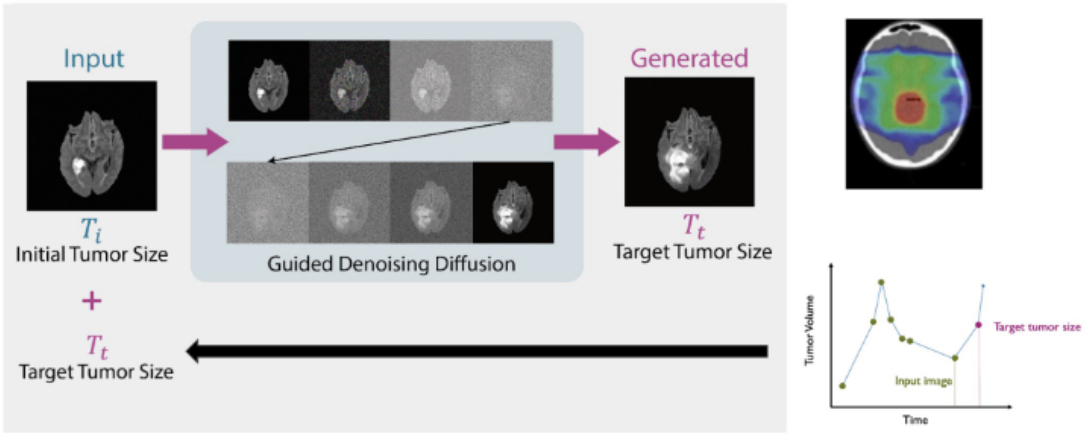

Figure 9: Conceptual overview regarding the integration of anatomical tumor growth modelling through spatial generative AI models and temporal mechanistic models of tumor growth.

5 Outcome of the Meeting

We successfully fulfilled the intended mission of bridging the gap between data science and mathematical modelling communities in oncology aiming to initiate new collaborations across disciplines. The event brought together a diverse and interdisciplinary group of participants, including data scientists, mathematicians, computational biologists, and clinical researchers, who actively participated throughout the workshop and beyond, leading to a variety of outputs generated as a result.

5.1 Participants' own perception

We initially assessed the participants' perception of the outcome of the workshop through a survey in the closing session. Here, almost half of the participants indicated that they started new collaborations during the workshop, aiming for future grant applications and joint projects and also indicated that people they met, they likely would not have made contact with outside the workshop. Most importantly, we also obtained very positive feedback regarding the overall format of the workshop and its outcomes as illustrated in selected participants' quotes (anonymous submissions via centimeter):

- "Very engaged and caring organizers, very respectful and productive atmosphere."
- "I loved the amount of time available for working in groups. The research talks were a great length."
- "It was lovely and very educational, a great balance of both fields. Thank you for a fantastic week!"
- "I met some wonderful people and started a new collaboration. The venue was wonderful."

5.2 Continued Collaborations and Lasting Impact

One of the most tangible outcomes of the workshop has been the continued collaboration among the multidisciplinary teams that were formed during the event. Participants were grouped into project teams based on complementary expertise and interest in specific subdomains of mathematical oncology and - given the timeframe - initiated but largely did not complete their anticipated projects during the workshop. A key objective for the organizers was to motivate new collaborations that would surpass the workshop. Encouragingly, the collaborative momentum has been sustained since January. A dedicated follow-up call in March demonstrated that all teams remain active, continuing their work with regular virtual meetings, collaborative code development, and shared document platforms. Teams have refined their models and expanded the scope

of their analyses, incorporating additional data modalities and theoretical frameworks introduced during the workshop. Some groups have even organized follow-up visits within Europe to continue their collaborative journey. We also noted a number of exchanged requests regarding invitations for international (virtual) seminar presentations, to bring the momentum of the workshop to the associated teams back home as well.

5.3 Scholarly Dissemination and Anticipated Contributions

A major aim of the workshop was to catalyze the production of tangible scholarly outputs. Several teams are currently preparing submissions to leading conferences and journals in the fields of mathematical and computational oncology, machine learning, and applied mathematics. At least three groups have indicated plans to submit abstracts to upcoming conferences such as the SIAM Conference on the Life Sciences, MICCAI, and ECCB, with draft manuscripts in progress for journals including PLOS Computational Biology and Journal of Computational Physics.

In addition to traditional academic outputs, each team has contributed a blog post to an established blog series dedicated to Mathematical Oncology, aimed at enhancing public engagement and cross-disciplinary visibility of the concept of mechanistic learning, here predominantly in the mechanistic modelling community. These blog posts distill complex technical advances into accessible narratives for a broader scientific and clinical audience. Topics include the integration of patient-specific -omics data into mechanistic learning models, interpretability in biology-informed neural networks, and the challenges of hybrid model validation - in line with the groups' activities. Samples of completed blog posts can be found at the following links:

- https://mathematical-oncology.org/blog/learning-dynamics-neural-nets.html
- https://mathematical-oncology.org/blog/machine-learning-for-model-parameterization.html, and
- https://mathematical-oncology.org/blog/insite.html.

The other submissions are currently being processed by the blog team and will be published in due course.

5.4 Conclusion and Future Outlook

The workshop has not only advanced the theoretical and practical frontiers of mechanistic learning but has also laid the foundation for long-term interdisciplinary collaboration. Participants have expressed strong interest in formalizing their group efforts into consortia for future joint funding applications. The success of the post-workshop follow-up call underscores the depth of engagement and the potential for sustainable scientific partnerships. Plans are underway to host follow-up symposia and webinars to continue community-building and share emerging results. This workshop has clearly demonstrated the value of integrating machine learning and mechanistic modelling in oncology, offering a replicable framework for interdisciplinary scientific collaboration. We greatly appreciate the support from the Banff International Research Station to make this possible and host our workshop!

References

[1] K. Willcox, O. Ghattas, and J. B. Freund, The imperative for physics-based modeling and inversion in machine learning, *Nature Computational Science* **1** (2021), 168–176.

[2] M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *Journal of Computational Physics* **378** (2019), 686–707.

[3] T. E. Yankeelov, J. C. Atuegwu, H. R. Hormuth II, et al., Clinically relevant modeling of tumor growth and treatment response, *Science Translational Medicine* **5**.187 (2013), 187ps9.

[4] R. Rockne, A. Hawkins-Daarud, K. R. Swanson, and P. Macklin, The 2019 mathematical oncology roadmap, *Physical Biology* **16**.4 (2019), 041005.

[5] H. B. Frieboes, F. Jin, Y. Chuang, et al., Integrated computational modeling of solid tumor progression and response to therapy, *Journal of Theoretical Biology* **264**.4 (2010), 1254–1270.

[6] A. Esteva, B. Kuprel, R. A. Novoa, et al., Dermatologist-level classification of skin cancer with deep neural networks, *Nature* **542** (2017), 115–118.

[7] C. Rudin, Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, *Nature Machine Intelligence* **1** (2019), 206–215.

[8] J. Metzcar, J. Nardini, S. Brüningk, and K. Wilkie, Mechanistic Learning in Mathematical Oncology: A Conceptual Landscape, in preparation (2024).

[9] S. Gu and W. Ng, Accelerated tumor growth simulations using deep learning surrogate models, *Computers in Biology and Medicine* **159** (2023), 106850.

[10] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proceedings of the National Academy of Sciences USA* **113**.15 (2016), 3932–3937.

[11] A. Kielland, M. J. Buckley, T. C. Yankeelov, and H. R. Hormuth II, Data-driven residual modeling for improved response prediction in glioblastoma, *Medical Image Analysis* **87** (2023), 102822.

[12] G. Currie, D. D. Hawk, A. M. Brady, et al., Digital twins for precision oncology: the next step in radiotherapy personalization, *The Lancet Oncology* **21**.12 (2020), e555–e561.

[13] P. Gerlee, R. Rockne, and K. R. Swanson, Predictive modeling in precision oncology, *Nature Reviews Cancer* **21** (2021), 270–282.

[14] J. Hestness, S. Narang, N. Ardalani, et al., Deep learning scaling is predictable, empirically, *arXiv preprint* arXiv:1712.00409 (2017).

[15] L. H. Gilpin, D. Bau, B. Z. Yuan, et al., Explaining explanations: An overview of interpretability of machine learning, in *Proceedings of the 2018 ICML Workshop on Human Interpretability in Machine Learning (WHI)*, (2018).

[16] M. Baker, Reproducibility crisis: Blame it on the antibodies, *Nature* **521** (2015), 274–276.

[17] L. Palmer and K. Willcox, Toward open, team science in engineering research, *Science Advances* **7**.10 (2021), eabe0133.