

Preface

Scientific discovery has long depended on human creativity, intuition, and expertise: the formulation of hypotheses, the design of experiments, the interpretation of evidence, and the construction of theories. In recent decades, however, advances in artificial intelligence (AI), machine learning, robotics, data science, and knowledge engineering have begun to reshape this landscape. AI for science, once a relatively specialized area, has become a major interdisciplinary field with growing scientific and societal significance. Computational systems now support—and in some settings partially automate—activities across the scientific process, from the analysis of complex data to model construction, experiment selection, and laboratory work.

The use of computation in scientific discovery has a long history, encompassing work in automated discovery, machine learning, knowledge representation, scientific modelling, cognitive science, and the philosophy of science. These traditions have often developed with different assumptions, representations, methods, and criteria of success. The AAAI 2023 Spring Symposium on Computational Approaches to Scientific Discovery was conceived as a forum in which these communities could meet, compare approaches, and identify shared problems. Its scope extended beyond any single technique: the central concern was how computational methods can contribute to the generation, evaluation, and refinement of scientific knowledge.

This concern has become increasingly urgent. The growth of scientific data, advances in computational power, and the expanding capabilities of automated experimentation have created opportunities for systems that reason over data and background knowledge, generate hypotheses, design informative experiments, and iteratively revise scientific models. Large language models and other foundation models add new possibilities for literature synthesis, knowledge acquisition, and machine-assisted reasoning. Scientific discovery goes beyond data processing. It involves uncertainty, abstraction, explanation, causal reasoning, and integration of heterogeneous forms of evidence and knowledge.

Successful use of computational approaches in science therefore requires more than improvements in predictive accuracy. It also calls for advances in representation and reasoning, interpretability, validation, reproducibility, and the organization of scientific workflows. As AI systems take more active roles in research, questions of trust, transparency, responsible use, and the division of labour between scientists and machines become increasingly important. Progress depends on dialogue across disciplines and on systems that combine data-driven learning with explicit scientific knowledge, experimental interaction, and forms of reasoning appropriate to scientific inquiry.

This volume grew out of the AAAI 2023 Spring Symposium on Computational Approaches to Scientific Discovery, held at the San Francisco Airport in Burlingame, California. The symposium brought together researchers from

artificial intelligence, machine learning, knowledge representation, robotics, laboratory automation, scientific modelling, cognitive science, philosophy of science, and a range of application domains. Its aim was not only to present recent results, but also to strengthen connections among communities pursuing computational scientific discovery through different paradigms and intellectual traditions.

The symposium attracted considerable interest, receiving 57 submissions. Fewer than half of these—28—were selected for oral presentation. Fourteen chapters in this volume grew out of these contributions and were substantially extended and revised for publication. The volume also contains two invited chapters that complement the symposium contributions. Every chapter underwent single-blind peer review by at least two referees, whose detailed comments and suggestions helped the authors improve their contributions.

The contributions collected here thus develop the conversation that started at the symposium in substantially extended form. Together, the sixteen chapters examine the foundations, methods, systems, and applications of computational scientific discovery. The intended readership includes researchers, postgraduate students, practitioners, and policymakers concerned with the development and responsible use of AI-enabled science.

The volume is organized into four thematic parts. Part I, “Automating Scientific Discovery,” considers the knowledge representations, conceptual foundations, inferential processes, experimental cycles, and robotic systems required to automate or support scientific inquiry. It moves from the representation and discovery of scientific concepts to integrated data- and knowledge-driven reasoning, closed-loop experimentation, and physically embodied robot scientists.

Part II, “Equation Discovery and Symbolic Regression,” focuses on methods for discovering interpretable mathematical expressions and symbolic relationships. The contributions present complementary approaches to this shared problem, including genetic programming, reinforcement learning, neural-guided search, grammar-based and neural generative models, and the automated generation of mathematical conjectures. Collectively, they show how scientific knowledge and structural constraints can guide search through large spaces of candidate expressions.

Part III, “Discovering Models of Dynamical Systems,” turns to explanatory and reduced-order representations of dynamic phenomena. It addresses quantitative process models, structural transformations of nonlinear ordinary differential equations, latent-space dynamics, progressive learning for physical systems, and process-based environmental modelling. Across these approaches, the emphasis is not merely on prediction, but on discovering models whose variables, mechanisms, and mathematical structure support scientific understanding.

The final chapter, “Artificial Intelligence for Science,” forming Part IV on its own, brings these three strands together and places them in the broader scientific, institutional, and policy landscape. Detailed accounts of the individual chapters are provided separately in the front matter, in the “Overview of the Contributions,” allowing this preface to concentrate on the origins, purpose, and intellectual organization of the book.

The chapters repeatedly demonstrate that computational discovery is most powerful when data-driven and knowledge-driven methods reinforce one another. Domain constraints, formal grammars, ontologies, process libraries, causal assumptions, and mathematical structures provide indispensable guidance through otherwise vast hypothesis spaces. They also reveal the growing importance of closed-loop systems in which hypotheses are generated, tested, and revised, all through interaction with data, simulations, or physical experiments. Such systems extend computational science beyond passive analysis toward active participation in inquiry.

Important challenges remain: scaling methods to complex systems; learning from sparse, noisy, or heterogeneous observations; discovering variables and processes not specified in advance; integrating knowledge across spatial and temporal scales; and evaluating the novelty, plausibility, explanatory power, and scientific value of machine-generated hypotheses. Addressing these challenges will require continued collaboration among computer scientists, domain scientists, engineers, philosophers, and experts in scientific practice.

AI systems are unlikely simply to replace scientists. Their more consequential role may be as partners that expand the scope of inquiry, accelerate the generation and testing of hypotheses, and make possible forms of experimentation and analysis that would otherwise be impractical. By presenting complementary perspectives on automated discovery, equation discovery, and modelling dynamical systems, this volume offers both an account of current progress and a foundation for further research into how intelligent systems can contribute responsibly and creatively to science.

We are grateful to everyone who contributed to the Symposium on Computational Approaches to Scientific Discovery and to the volume that grew from it. We thank all the authors who submitted work for consideration at the symposium, and everyone who attended and took part in the discussions there, particularly those who presented their research. We are especially grateful to the authors of all sixteen chapters collected in this volume, whose sustained effort turned symposium contributions and newly invited work alike into the book before you. We also thank the reviewers, whose careful reading and constructive comments substantially improved every chapter. We would next like to thank the Springer staff for their constant support and encouragement in bringing this volume to publication. Finally, we thank our institutions and funders for their support of this work. In particular, S. Džeroski gratefully acknowledges funding through the grants GC-0001 (AI for Science) and P2-0103 (Knowledge Technologies) by ARIS, the Slovenian Research and Innovation Agency and the HE Grant SLAIF 101254461 (funded by EuroHPC JU and Slovenian Ministry of Education, Science and Youth).

August 2026

Sašo Džeroski, Youngsoo Choi, J. Nathan Kutz
Pat Langley, Larisa Soldatova

Overview of the Contributions

The volume is organized into four parts that move from the general requirements and architectures of automated discovery, through methods for discovering equations, to approaches for constructing models of dynamical systems. A final synthesis chapter places these developments within the broader landscape of artificial intelligence for science.

Part I — Automating Scientific Discovery

The chapters in Part I examine the conceptual, representational, inferential, and experimental components needed to automate scientific discovery. They progress from the engineering of machine-processable scientific knowledge, through the discovery of scientific concepts and the integration of data with knowledge, to closed-loop experimentation in behavioral science and fully embodied robot scientists.

The chapter 'Knowledge Engineering for Automated Scientific Discovery' by Larisa Soldatova examines automated discovery systems through the lens of data and knowledge representation, with particular emphasis on the design principles that enable intelligent systems to operate as autonomous or semi-autonomous scientific agents. It further considers the growing role of Large Language Models as potential domain experts and reasoning components within discovery pipelines. A major focus of the chapter is the development and application of ontologies and structured knowledge resources for representing both scientific knowledge and the scientific process itself. The discussion includes ontologies designed for experimental investigations, hypotheses, and computational workflows. These approaches are illustrated through a case study of robot scientists, i.e. 'Adam', 'Eve', and 'Genesis', highlighting how domain knowledge, procedural knowledge, and experimental outputs can be encoded to support autonomous reasoning and experimentation. The chapter argues that advances in AI-driven discovery and laboratory automation are intensifying the challenges associated with large-scale scientific data and dynamic background knowledge. Addressing these challenges requires a new generation of integrated data and knowledge infrastructures capable of supporting software agents, adaptive reasoning, and closed-loop scientific discovery. Ultimately, holistic representations of scientific investigations may not only improve the reliability and scalability of automated discovery systems, but also contribute to a deeper computational understanding of the nature of scientific inquiry itself.

The chapter 'The Relevance of Natural Kinds to the Computational Discovery of Novel Scientific Variables' by Benjamin Jantzen discusses the role of natural kinds and dynamical similarity in the computational discovery of scientific knowledge, focusing on two closely related problems: determining whether different systems share a common underlying causal structure, and identifying

novel scientific variables that expand the dimensions through which physical phenomena can be described. The chapter argues that a central goal of scientific inquiry is the identification of “projectible” classes of systems—groups of phenomena for which observations and laws discovered in one instance can be generalized to others. While fundamental physics seeks universal laws, many scientific disciplines instead rely on identifying more restricted dynamical kinds governed by shared causal principles, such as systems studied in chemistry, thermodynamics, or electronic circuits; within this context, the discovery of novel variables becomes essential for improving explanation, prediction, and control by increasing the expressive power of scientific models and enriching the state spaces used to describe physical systems. The chapter presents a computational framework that connects natural kind identification with automated variable discovery, reformulating the qualitative problem of dynamical kind identification as a quantitative computational task and developing methods capable of operating even when the governing laws of the systems are unknown. These methods are integrated into an iterative discovery cycle in which computational assessment of dynamical similarity guides the introduction of novel variables, followed by experiments that reveal the causal significance of these variables. Experimental results from studies of water flow through granular media are used to illustrate this discovery process in practice.

The chapter ‘Data-Driven and Knowledge-Based Approach to Scientific Discovery’ by Mitsuhiro Odaka, Morgan Magnin, and Katsumi Inoue presents a Data-Driven and Knowledge-Based (DD-KB) approach to scientific discovery that aims to bridge the growing gap between large-scale data accumulation and the generation of scientific knowledge. The proposed framework integrates data-centric computational methods with explicit knowledge representation and logical reasoning within a closed discovery loop involving deduction, induction, and abduction. Central to the approach is the continuous transformation between data and knowledge through data-to-knowledge conversion, knowledge-to-data generation, and embedding symbolic knowledge into continuous vector spaces to accelerate reasoning processes, supporting the construction of interpretable and parsimonious scientific models, including logic programs, equations, and causal networks. The chapter illustrates the framework through two case studies focused on understanding the biological mechanisms of COVID-19 caused by SARS-CoV-2 infection. The first study investigated viral dynamics at a macroscopic scale using differential equation models, sensitivity analysis, stability analysis, and system identification techniques applied to viral load data from hospitalized patients, generating the hypothesis that direct cell-to-cell virus transmission may be associated with disease severity. The second study explored this hypothesis at the microscopic scale by examining signaling pathways associated with intercellular adhesion molecule 1 (ICAM-1), combining single-cell omics data mining, gene network inference, and validation against biological knowledge bases; the resulting analyses reproduced known pathways while also proposing previously unknown ICAM1-associated signaling mechanisms potentially linked to SARS-CoV-2 cell-to-cell transmission. The chapter concludes by discussing the broader

implications of the DD-KB framework for automated scientific discovery, highlighting both its potential and its current limitations related to scalability, data quality, hypothesis validation, and the continuing importance of human expertise in experimental design and interpretation.

The chapter 'Automated Scientific Discovery of Mind and Behavior' by Sebastian Musslick, Younes Strittmatter, and Marina Dubova explores the emerging paradigm of automated scientific discovery in the behavioral and cognitive sciences, focusing on closed-loop discovery systems that integrate experimental design, data collection, and model inference into an integrated computational process. The chapter formalizes behavioral research as an iterative computational cycle in which automated systems can actively participate in generating experiments, collecting behavioral data, and inferring explanatory models, emphasizing the multiple epistemic goals of behavioral models, including description, prediction, explanation, and behavioral control. The chapter presents AutoRA, an open-source framework designed to implement closed-loop behavioral research workflows, integrating automated experimental design, data collection, and model inference into a dynamic discovery cycle and serving as an initial proof of principle for fully automated behavioral scientific discovery. The authors examine major challenges facing automated scientific discovery in the behavioral sciences, including the representation of experiments and theories, evaluation of scientific progress, integration across diverse modeling frameworks, and the need for structured, machine-readable scientific knowledge representations such as knowledge graphs. The chapter concludes by arguing that automated discovery systems may fundamentally reshape behavioral and cognitive science by enabling more transparent, scalable, and computationally integrated approaches to understanding mind and behavior, while still requiring meaningful interaction between human researchers and automated systems.

The chapter 'The Use of AI-Robotic Systems for Scientific Discovery' by Gower et al. examines the use of AI-driven robotic systems, or "robot scientists," for automating scientific discovery through the integration of artificial intelligence, laboratory robotics, and computational reasoning. The authors present robot scientists as systems capable not only of generating hypotheses and models from data, but also of autonomously designing, executing, and evaluating real-world experiments within a closed scientific discovery loop, exploring the philosophical foundations of robot scientists—theories, models, logical and statistical inference, parsimony, and theory choice. It further argues that the scientific method shares a close analogy with active learning rather than traditional reinforcement learning, since scientific discovery involves selectively performing informative experiments to reduce uncertainty and improve explanatory models. The chapter illustrates these ideas through the evolution of robot scientists—Adam, Eve, and the next-generation platform Genesis. Adam demonstrated autonomous hypothesis generation and testing in yeast metabolism research using logic programming and automated experimentation, while Eve automated aspects of drug discovery through quantitative structure–activity relationship modeling and iterative compound testing. Genesis extends these capabilities through a large-scale

microfluidic experimental platform containing 1000 computer-controlled micro-bioreactors integrated with high-throughput measurement technologies such as mass spectrometry and RNA sequencing. The chapter also introduces LGEM+, a first-order logic framework grounded in genome-scale metabolic models and automated theorem proving, which supports hypothesis generation and reasoning within Genesis, and concludes by discussing future directions toward fully integrated closed-loop experimentation, where computational models, laboratory robotics, controlled vocabularies, and machine learning systems work together to support scalable and autonomous scientific discovery.

Part II — Equation Discovery and Symbolic Regression

Part II turns to the automated discovery of interpretable mathematical expressions. Its chapters present complementary ways to organize and guide search through vast spaces of candidate equations. The chapters are arranged from scientifically inspired genetic programming, through reinforcement learning and neural-guided search, to grammar-based and neural generative models, closing with the automated generation of mathematical conjectures in mathematics.

The chapter ‘Expediting Symbolic Regression for Science Using Scientific Approaches’ by Md Nasim, Yexiang Xue, and Nan Jiang provides an overview of computational approaches that accelerate scientific discovery through symbolic regression and scientifically inspired reasoning strategies, reviewing several influential systems, including the pioneering discovery systems BACON and LA-GRANGE, as well as the autonomous “robot scientists” platforms Adam and Eve. The chapter focuses particularly on symbolic regression, a core challenge in AI-driven scientific discovery concerned with uncovering interpretable mathematical equations directly from experimental data. To address the difficulty of discovering equations automatically in systems with many interacting variables, the chapter introduces Control Variable Genetic Programming (CVGP), which incrementally constructs complex equations by first discovering simpler reduced-form relationships under controlled experimental conditions, then progressively introducing new variables through additional control-variable experiments. In addition, the chapter discusses broader principles underlying scientific approaches to discovery, including problem specification, hypothesis formulation, experimental testing, and data analysis.

The chapter ‘Deep Symbolic Optimization: Reinforcement Learning for Symbolic Mathematics’ by Conor Hayes et al., aka. the DSO (Deep Symbolic Optimization) Team, presents DSO—a computational framework that combines deep learning, reinforcement learning, and symbolic reasoning to address challenging symbolic optimization problems in scientific discovery. In the DSO framework, symbolic optimization is formulated as a sequential decision-making problem in which an autoregressive neural network generates candidate symbolic expressions token by token, while reinforcement learning guides the search toward expressions with high fitness or explanatory value, incorporating domain knowledge through in-situ priors and constraints. The chapter provides an overview of

the modular DSO architecture and describes how reinforcement learning, evolutionary computation, imitation learning, and large-scale pre-training can be integrated into a unified symbolic optimization framework (uDSO), demonstrating state-of-the-art performance in both predictive accuracy and interpretability. In addition to symbolic regression, the chapter discusses broader applications of DSO in areas such as reinforcement learning policy discovery, power converter design, antibody optimization, and medical treatment optimization, and outlines future directions involving multi-objective optimization and the integration of large language models.

The chapter ‘Neural-Guided Equation Discovery’ by Brugger et al. examines neural-guided approaches to equation discovery, focusing on how deep learning techniques can be integrated with symbolic search methods to efficiently identify mathematical expressions from data. The authors present MGMT (Multi-Task Grammar-Guided Monte-Carlo Tree Search for Equation Discovery), a modular framework that combines neural-guided Monte-Carlo Tree Search (MCTS), supervised and reinforcement learning, and context-free grammars to guide the search for candidate equations. A central theme of the chapter is the investigation of how neural networks should best support symbolic search, including whether equation generation should rely on grammar rules or token-based representations, how tabular data should be embedded, and whether supervised or reinforcement learning offers more effective training strategies. The chapter presents extensive experimental evaluations of MGMT across benchmark symbolic regression problems and a real-world application involving prediction of droplet friction forces, comparing recurrent, convolutional, and transformer-based architectures as well as multiple adaptations of MCTS such as risk-seeking MCTS and AmEx-MCTS. Their results show that neural-guided search substantially improves search efficiency and that supervised learning generally outperforms reinforcement learning for guiding equation discovery.

The chapter ‘Generative Models for Equation Discovery’ by Mežnar et al. is devoted to the role of generative models in equation discovery, focusing on methods that automatically generate candidate mathematical expressions within the “generate-and-test” paradigm. The chapter explores two major classes of generative approaches: manually crafted generators based on probabilistic context-free grammars (PCFGs), and trained neural generative models based on variational autoencoders (VAEs). The chapter presents ProGED as a representative grammar-based framework that generates candidate equations through stochastic sampling guided by probabilistic grammars and Bayesian updating mechanisms, and EDHiE (Equation Discovery with Hierarchical variational autoEncoders), a framework combining hierarchical variational autoencoders with evolutionary search techniques for discovering mathematical expressions. Comparative experiments on reconstructing Coulomb’s law from data show that both grammar-based and neural generative approaches can successfully recover correct physical equations. The chapter concludes by advocating a generator-centric perspective on equation discovery, emphasizing that future progress depends on developing generative systems capable of integrating prior scientific knowledge,

adapting to data, and imposing meaningful structure on the enormous space of possible symbolic expressions.

The chapter 'Ramanujan Machine—Generating Conjectures on Fundamental Constants' by Beit-Halachmi et al. explores how algorithms can assist in one of the most creative aspects of mathematics: the discovery of new formulas and conjectures. Focusing on fundamental mathematical constants such as π , e , and values of the Riemann zeta function, the chapter presents the Ramanujan Machine project as a landmark effort in experimental mathematics that automates the generation of mathematical conjectures. The chapter surveys several algorithmic approaches, including the Meet-in-the-Middle with Regular Formulas (MITM-RF) algorithm, gradient-based Descent and Repel methods, and the ESMA algorithm, which combines sequence extraction with pattern detection. Beyond generating individual formulas, the chapter highlights how automated conjecture generation can lead to the discovery of entirely new mathematical structures and concepts, including the Factorial Reduction property in families of continued fractions and the Conserving Matrix Field (CMF) framework, which appears to unify broad classes of conjectured formulas and generalizes aspects of Apéry's famous proof of the irrationality of $\zeta(3)$. The authors position this work within the emerging field of Automated Conjecture Generation, complementing the more established area of automated theorem proving. The chapter thus presents a compelling vision of human–AI collaboration in mathematics, where computational systems act as partners in the discovery process, generating insights that inspire new mathematical theories and proofs.

Part III — Discovering Models of Dynamical Systems

The chapters in Part III move beyond isolated symbolic expressions to the structure, construction, reduction, transfer, and application of models of dynamical systems. They begin with explanatory process models, then examine alternative polynomial and quadratic representations of nonlinear ordinary differential equations, latent-space models of high-dimensional physical systems, transferred reduced-order models, and automated process-based modelling in environmental science.

The chapter 'Computational Discovery of Quantitative Process Models' by Pat Langley provides a comprehensive review of two decades of research on the automated discovery of quantitative process models—scientific models that explain dynamic phenomena in terms of interacting processes rather than merely describing observed data. Langley argues that scientific understanding ultimately requires explanatory models that capture the underlying mechanisms responsible for observed behaviour, introducing a formal representation in which scientific processes are described through variables, rate expressions, and differential equations. The chapter traces the evolution of computational methods for inducing these models, from early systems such as IPM (Inductive Process Modeling), which relied on exhaustive search and gradient-based parameter estimation but suffered from scalability and reliability limitations, to increasingly sophis-

ticated systems—RPM and later SPM—that introduced heuristic search, linear regression-based parameter estimation, backward elimination, and sampling methods. Beyond reviewing these advances, Langley identifies several important future challenges, including handling unobserved variables, discovering entirely new processes, integrating experimentation with model induction, and extending the framework to spatio-temporal and multi-scale phenomena. The chapter positions quantitative process modelling as a distinctive branch of computational scientific discovery that combines data, domain knowledge, causal reasoning, and mechanistic explanations, offering a pathway toward more interpretable and scientifically meaningful AI systems for understanding complex dynamic systems.

The chapter ‘Discovering Polynomial and Quadratic Structure in Nonlinear Ordinary Differential Equations’ by Boris Kramer and Gleb Pogudin examines the discovery of polynomial and quadratic structure in nonlinear ordinary differential equations (ODEs), focusing on methods that transform complex nonlinear dynamical systems into polynomial or quadratic forms. Such transformations, known as polynomialization and quadratization, reveal hidden structural representations of dynamical systems by introducing new variables through a process often referred to as lifting. Building on mathematical results developed over more than a century, the authors review theoretical foundations showing that many nonpolynomial systems can be reformulated into polynomial—and often quadratic—representations. The chapter surveys algorithmic methods and software tools that automate these transformations, discussing challenges associated with finding optimal quadratizations, a problem known to be computationally difficult. To illustrate the practical relevance of these methods, the chapter presents applications involving a single-layer neural network and a phenomenological model of cell signaling, while also highlighting broader applications in areas including chemical reaction networks, combustion modeling, fluid dynamics, and additive manufacturing.

The chapter ‘A Comprehensive Review of Latent Space Dynamics Identification Algorithms for Intrusive and Non-Intrusive Reduced-Order-Modeling’ by Christophe Bonneville et al. presents a comprehensive review of Latent Space Dynamics Identification (LaSDI), a family of machine-learning-based reduced-order modeling methods designed to accelerate high-fidelity simulations of physical systems governed by partial differential equations. LaSDI combines nonlinear dimensionality reduction, typically through autoencoders, with data-driven identification of ordinary differential equations in a low-dimensional latent space. The chapter surveys both intrusive and non-intrusive variants of the LaSDI framework, including methods that incorporate thermodynamic structure, weak-form equation learning, active learning, Gaussian-process-based interpolation, and uncertainty quantification. Through examples involving Burgers’ equation, nonlinear heat conduction, and plasma physics, the chapter demonstrates that LaSDI-based methods can achieve high accuracy with substantial computational speedups, making them promising tools for scientific machine learning, digital twins, uncertainty quantification, design optimization, and decision-making workflows.

The chapter ‘Progressive Learning for Subsurface Physics’ by Kadeethum et al. presents progressive reduced-order modeling (p-ROM), a transfer-learning framework designed to improve data-driven reduced-order models for complex subsurface physics problems. The p-ROM framework enables a new “child” model to learn from multiple pretrained “parent” models, allowing knowledge to be transferred across different physics domains, parameter spaces, geometries, and dimensional settings while reducing the risk of catastrophic forgetting. The chapter examines the trade-offs involved in progressive learning, emphasizing that adding more parent models does not always improve accuracy and may increase architectural complexity and computational cost; to better understand when transferred knowledge is beneficial, the authors introduce layer attribution analysis as a diagnostic tool for quantifying how much each parent model contributes to the final prediction. Through examples involving porous media transport, hyperelastic material deformation, and transfer across two- and three-dimensional settings, the chapter demonstrates how p-ROM can reduce training data requirements while maintaining predictive accuracy.

The chapter ‘Process-Based Modelling of Natural and Urban Catchments’ by Radinja et al. presents a comprehensive framework for automated environmental modelling based on the Process-Based Modelling Tool (ProBMoT). Rather than relying on manually assembled models or purely data-driven black-box methods, the chapter proposes an equation-discovery approach that combines measured environmental data with structured domain knowledge encoded in reusable libraries of entities and processes; ProBMoT systematically generates candidate model structures, transforms them into systems of algebraic and differential equations, calibrates their parameters using optimization techniques such as differential evolution, and ranks the resulting models according to predictive performance. In the Quarteira River catchment in Portugal, ProBMoT automatically generated and calibrated multiple hydrological and water-quality models, with automated exploration of alternative formulations outperforming fixed traditional modelling approaches by selecting more appropriate combinations of evapotranspiration, sediment transport, and nutrient-loading equations. In an urban catchment case study in Ljubljana, Slovenia, ProBMoT was used to generate rainfall-runoff models and support the design of stormwater control measures under different hydrological scenarios. Across both domains, the use of automated process-based modelling reduces human bias, improves transparency, and supports informed environmental management decisions. Overall, the chapter positions ProBMoT as a powerful and extensible framework for integrating domain expertise, equation discovery, and automated reasoning in environmental modelling and sustainable catchment management.

Part IV — Artificial Intelligence for Science

The final part consists of a synthesis chapter that broadens the perspective from individual discovery methods and systems to the wider methodological and infrastructural requirements of artificial intelligence for science.

The chapter 'Artificial Intelligence for Science' by Sašo Džeroski surveys AI approaches designed to meet challenges distinctive to scientific research: models must often be explainable and interpretable, learn from limited labelled data, incorporate substantial prior knowledge, and support the open and reproducible representation and exchange of results. The chapter relates these requirements to explainable machine learning—using decision trees, rule sets, ensembles, and multi-target prediction as central examples—and to semi-supervised learning for settings in which labelled observations are scarce; it then considers foundation models and their adaptation to scientific tasks through transfer learning and fine-tuning, automated scientific modelling through the induction of interpretable differential-equation systems from time-series data, and semantic technologies and ontologies that enable open science. Applications in the life, environmental, and materials sciences illustrate how these methods operate in practice. The chapter concludes by turning from methods to enabling infrastructure, discussing AI factories in general and the Slovenian AI Factory (SLAIF) in particular as an example of the computational and institutional capacity needed to advance AI-enabled science.