

Knowledge-Guided Machine Learning: A Paradigm Shift in AI for Science

Anuj Karpatne¹, Xiaowei Jia², Vipin Kumar³

¹Virginia Tech

²University of Pittsburgh

³University of Minnesota

karpatne@vt.edu, xiaowei@pitt.edu, kumar001@umn.edu

Abstract

As advances in artificial intelligence (AI) and machine learning (ML) continue to transform commercial applications, the scientific community is increasingly eager to harness AI/ML's power to accelerate modeling and discovery. However, purely data-driven AI methods often lack interpretability, generalizability, and consistency with established scientific principles. Conversely, traditional process-based models embody deep scientific knowledge but suffer from limited scalability or incomplete representation of complex systems. Knowledge-guided machine learning (KGML) offers a promising path forward by integrating scientific knowledge with data-driven approaches to produce AI models that are robust, trustworthy, and capable of advancing both AI and science. This talk summarizes the foundations of KGML, outlines a taxonomy for organizing research efforts, and highlights emerging opportunities for broad scientific impact.

Introduction

Recent advances in AI and machine learning (ML) have led to remarkable progress in fields such as computer vision and natural language processing (Wang et al. 2023; Zhang et al. 2023). These successes are now inspiring the scientific community to adopt AI/ML as a tool for accelerating modeling and discovery. However, applying ML in scientific contexts presents unique challenges. For example, scientific data often suffer from sparse availability of labels from expensive observations and experiments. Also, AI solutions in science must respect established physical, biological, or chemical principles to be used as reliable aides to scientists for generating and validating scientific hypotheses (Krenn et al. 2022; AI4Science and Quantum 2023). While black-box ML models that are supervised solely by data excel at learning patterns given large volumes of data, they are prone to overfitting and producing predictions that violate basic scientific laws especially in out-of-distribution scenarios.

In contrast to using data-only methods, the conventional approach for modeling scientific systems is to use process-based or mechanistic models that are anchored in scientific theories and principles. However, process-based models suffer from several shortcomings limiting their adoption in complex real-world settings, e.g., due to imperfections in

model formulations (or modeling bias), incorrect choices of parameter values in equations, and high computational costs in running high-fidelity simulations.

Given the complementary strengths and limitations of purely data-driven AI and purely mechanistic scientific models, there is increasing interest in integrating scientific knowledge into AI/ML models in the rapidly growing field of *knowledge-guided machine learning (KGML)* (Karpatne et al. 2017; Willard et al. 2022; Karpatne, Kannan, and Kumar 2022; Karpatne, Jia, and Kumar 2024). Incorporating knowledge directly in the solution structure or the learning algorithm of AI models provides a mechanism to improve the scientific consistency and generalizability of AI solutions on out-of-distribution data while being parameter and label efficient. KGML provides a systematic path for aligning heterogeneous data sources in scientific applications with known laws and equations, producing AI models that can merge information across diverse representations and modalities of scientific systems. Also, by embedding scientific knowledge in the reasoning process of AI systems, KGML aims at improving the interpretability and trustworthiness of AI solutions to be used as credible aides in the process of scientific discovery.

In this talk, we present a brief overview of the research happening in KGML by introducing a taxonomy of KGML approaches. We then discuss some emerging opportunities and challenges at the frontiers of KGML research, followed by concluding remarks. This talk is an abridged summary of a longer review article on KGML (Karpatne, Jia, and Kumar 2024).

A Taxonomy of KGML Approaches

Research in KGML spans multiple dimensions that together define how knowledge and data-driven models can be integrated in the context of a diverse range of scientific applications. Figure 1 provides a schematic overview of these dimensions. We briefly describe each of these dimensions in the following.

Format of Scientific Knowledge

Scientific knowledge can be represented in diverse formats, ranging from partial differential equations (PDEs) that govern physical systems, to conservation laws, phenomenological rules, and qualitative heuristics. In some do-

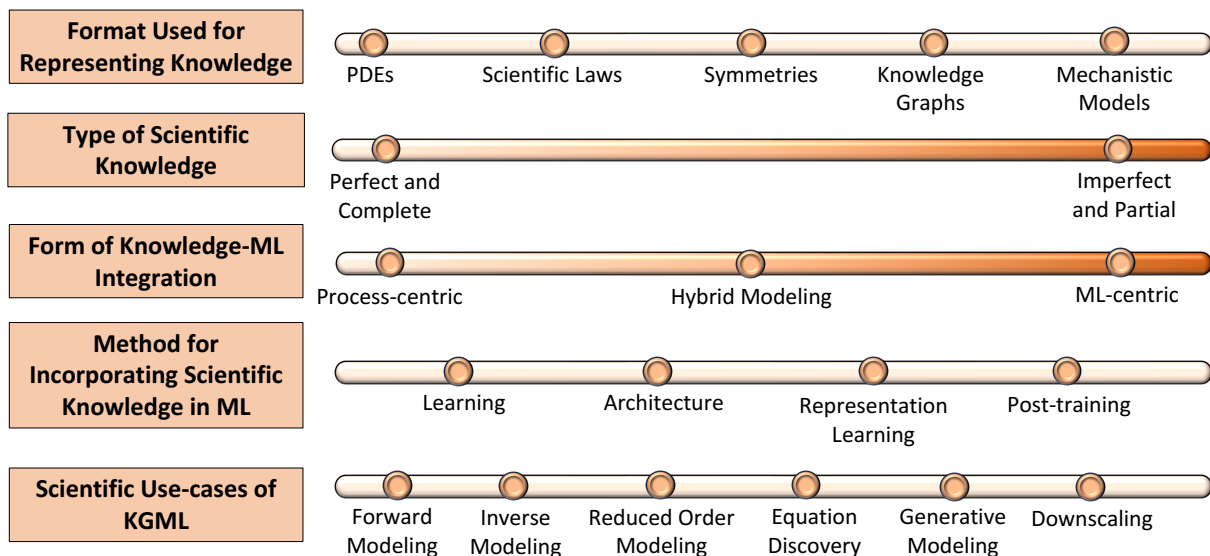


Figure 1: A multi-dimensional view of KGML research to characterize the diversity of research threads combining scientific knowledge with ML.

mains, knowledge is expressed as symmetries and invariances, while in others it is encoded in the form of knowledge graphs and ontologies. Knowledge may also be embodied in the form of process-based or mechanistic models. Each format presents distinct opportunities and challenges for integrating knowledge with ML.

Type of Scientific Knowledge

Along with the format, the fidelity and completeness of available knowledge also varies greatly across scientific applications. In some problems, knowledge can be assumed to be perfect and complete, e.g., when solving idealized PDEs where governing equations are fully specified but computationally expensive to solve numerically (Raissi, Perdikaris, and Karniadakis 2019). In such problems, KGML presents a promising alternative to build surrogates of numerical models that improve the speed and scalability of scientific simulations without comprising of the accuracy. However, there are many other scientific applications where knowledge is partial (or incomplete) and imperfect, with approximations in certain terms or gaps in mechanistic understanding (Jia et al. 2019). In such problems, the role of KGML models is to also improve the predictive accuracy of solutions with respect to ground-truth observations. There is thus a continuum of problems in science where knowledge varies from perfect and incomplete to imperfect and partial, motivating diverse KGML formulations.

Form of Knowledge-ML Integration

Another key dimension is how strongly knowledge and ML share control over the final modeling process. On one side of this spectrum, ML-centric approaches let data-driven methods dominate while using knowledge only as a guide or constraint to influence the ML model design or learning algo-

rithm. On the other hand, process-centric approaches rely primarily on mechanistic models to represent the mapping of inputs to outputs, while the role of ML models is limited to calibrating parameters, inferring intermediate quantities, and learning corrective terms in process-based equations. This dimension is again a continuum where KGML approaches can range from purely process-centric to ML-centric. Indeed, at the middle of this continuum sit hybrid modeling methods where both process-based and ML models play primary roles in mapping inputs to outputs.

Method for Incorporating Scientific Knowledge in ML

For any ML component used in a KGML setup, we can incorporate knowledge in ML frameworks in a variety of ways. One approach is modifying the learning algorithm, such as adding loss terms that penalize violations of known laws. Another is designing architectures that embed knowledge directly into the structure of ML models such as neural network architectures. A third approach is representation learning, where feature spaces learned by ML models are grounded in scientific principles. A fourth approach is post-training, where pre-trained AI models can be fine-tuned or guided during inference using scientific knowledge. Note that these methods are not mutually disjoint and can also be combined together.

Scientific Use-cases in KGML

Finally, KGML methods serve a variety of use-cases in scientific applications. A common use-case is forward modeling, where ML guided by knowledge learns mappings from drivers to response variables. Another use-case is inverse modeling, where KGML recovers hidden parameters or states of mechanistic models from data, aiding calibra-

tion and system understanding (Tsai et al. 2021). KGML can also be used for reduced order modeling, where the goal is to learn a lower-dimensional representation of a more complex set of scientific equations. A fourth use-case is in discovery of scientific equations automatically from data. A fifth use-case involves generative modeling, where KGML approaches are used to build digital twins of scientific systems capable of producing realistic synthetic data. A sixth use-case is downscaling, where coarse-resolution data are transformed into fine-scale predictions, reducing the computational cost of high-resolution simulations. Note that this is not exhaustive list; there are several other use-cases in science where AI/ML methods can be used, providing another dimension to differentiate KGML formulations.

Opportunities and Challenges

There are several opportunities in advancing the frontiers of KGML for accelerating the field of AI for Science. One key opportunity is improving out-of-distribution generalization especially when the training data is limited in size and variety. By grounding AI models in scientific principles that can be assumed to universally hold true across a large variety of testing scenarios not seen in the training data, KGML enables extrapolation beyond the training data, addressing a central limitation of black-box (or data-only) ML models. Another opportunity lies in the advancement of scientific knowledge—by combining mechanistic reasoning with data-driven flexibility, KGML systems can reveal novel relationships and gaps in existing knowledge in the form of scientific hypotheses, guiding new experiments to even validate such hypotheses.

In addition, KGML opens avenues for interdisciplinary collaboration. Scientists from diverse domains can now share modeling strategies grounded in both data and knowledge, creating opportunities for cross-pollination between fields such as physics, biology, chemistry, and engineering. Building such bridges across AI and scientific and engineering disciplines requires shared benchmarks and community-driven platforms.

At the same time, there are several challenges that remain to be addressed in the field of KGML. While some formats of knowledge representation has received considerable attention in the KGML community (e.g., solving PDEs), there are many other formalizations of knowledge that require similar attention (e.g., problems where scientific knowledge is imperfect or incomplete in the form of approximate hypotheses). Furthermore, we need standardized benchmark datasets and evaluation protocols to quantify our progress in KGML for improving out-of-distribution generalizability and interpretability across a wide range of scientific and engineering applications. There is also an urgent need for training a new generation of researchers fluent in both domain sciences and AI, capable of advancing this interdisciplinary field forward. Finally, we need to account for ethical considerations around bias, transparency, sustainability, and equitable access to AI infrastructure in scientific applications.

Conclusions

Knowledge-guided machine learning represents a paradigm shift in AI for science. By combining the explanatory power of scientific knowledge with the adaptability of ML, KGML offers a path toward models that are accurate, interpretable, and grounded in scientific reality. KGML can catalyze a new wave of AI-driven discoveries, creating a virtuous cycle where advances in AI propel science forward, and the demands of science inspire new innovations in the science of AI.

References

- AI4Science, M. R.; and Quantum, M. A. 2023. The Impact of Large Language Models on Scientific Discovery: a Preliminary Study using GPT-4. *arXiv preprint arXiv:2311.07361*.
- Jia, X.; Willard, J.; Karpatne, A.; Read, J.; Zwart, J.; Steinbach, M.; and Kumar, V. 2019. Physics guided RNNs for modeling dynamical systems: A case study in simulating lake temperature profiles. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, 558–566. SIAM.
- Karpatne, A.; Jia, X.; and Kumar, V. 2024. Knowledge-guided machine learning: Current trends and future prospects. *arXiv preprint arXiv:2403.15989*.
- Karpatne, A.; Kannan, R.; and Kumar, V. 2022. *Knowledge Guided Machine Learning: Accelerating Discovery Using Scientific Knowledge and Data*. CRC Press.
- Karpatne, A.; Watkins, W.; Read, J.; and Kumar, V. 2017. Physics-guided Neural Networks (PGNN): An Application in Lake Temperature Modeling. *arXiv preprint arXiv:1710.11431*.
- Krenn, M.; Pollice, R.; Guo, S. Y.; Aldeghi, M.; Cervera-Lierta, A.; Friederich, P.; dos Passos Gomes, G.; Häse, F.; Jinich, A.; Nigam, A.; et al. 2022. On scientific understanding with artificial intelligence. *Nature Reviews Physics*, 4(12): 761–769.
- Raïssi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378: 686–707.
- Tsai, W.-P.; Feng, D.; Pan, M.; Beck, H.; Lawson, K.; Yang, Y.; Liu, J.; and Shen, C. 2021. From calibration to parameter learning: Harnessing the scaling effects of big data in geoscientific modeling. *Nature communications*, 12(1): 5988.
- Wang, H.; Fu, T.; Du, Y.; Gao, W.; Huang, K.; Liu, Z.; Chandak, P.; Liu, S.; Van Katwyk, P.; Deac, A.; et al. 2023. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972): 47–60.
- Willard, J.; Jia, X.; Xu, S.; Steinbach, M.; and Kumar, V. 2022. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Computing Surveys*, 55(4): 1–37.
- Zhang, X.; Wang, L.; Helwig, J.; Luo, Y.; Fu, C.; Xie, Y.; Liu, M.; Lin, Y.; Xu, Z.; Yan, K.; et al. 2023. Artificial intelligence for science in quantum, atomistic, and continuum systems. *arXiv preprint arXiv:2307.08423*.