


# Optimizing the synergy between physics and machine learning

Although the initial inspiration of neural networks came from biology, insights from physics have helped neural networks to become usable. New connections between physics and machine learning produce powerful computational methods.

**M**achine learning and physics have long-standing strong links. An important connection was forged in 1982 by John Hopfield, as he considered the analogy between a physical system that consisted of interacting particles in which emergent phenomena such as magnetism arise, and a network of interacting neurons with spontaneous computational properties. The Hopfield network is a forerunner of the recurrent neural network, which has its use in a wide range of machine learning applications that involve temporal, dynamic features.

From a different perspective, theoretical physics is expected to help with a foundational understanding of machine learning. A [classic paper](#) from 1984 by L. G. Valiant set the tone, describing a rigorous statistical theory of learning. But with the rise of deep learning since the 2010s, further questions have emerged about the surprising, unreasonably good performance and generalization capabilities of deep neural networks. In a recent [Comment](#) in *Nature Physics*, Lenka Zdeborová calls for renewed efforts to tackle such questions with physics-inspired approaches, pointing to physicists' experience with tackling observations from a large number and varied range of experiments by searching for models that can capture the essence of a problem, ignoring many of the details, and testing it with analytical investigations. A case in point is the highly successful Ising model of magnetism, which does not include any of the quantum mechanical details of the magnetic interactions or material properties but explains many different types of experimental phenomena. In fact, the Ising model has proved useful in machine learning too, as Hopfield's network can be cast as an Ising model of a neural network. Efforts in this direction can help to develop a [theoretical understanding of deep learning](#) and why it works so well.

From the point of view of practical applications, there are many opportunities to put state-of-the-art machine learning models to good use, tackling problems



A simplified, very rough energy landscape, for protein folding. Source: Kuhlman & Bradley. *Nat. Rev. Mol. Cell Biol.* **20**, 681–697 (2019). Credit: Springer Nature Ltd

in physical sciences whenever large amounts of data are involved, such as in [fluid mechanics](#), [high-energy physics](#) or [weather forecasting](#). An active topic is the use of machine learning models to learn the mathematical rules, in the form of partial differential equations, that underlie complex dynamic phenomena such as turbulence. In a recent [Review article](#) in *Nature Reviews Physics*, George Karniadakis and colleagues discuss various ways in which physics can be embedded in such approaches to tackle complex dynamic problems, from earthquake predictions to molecular dynamics.

Another promising connection that currently sees [much activity](#) is tackling optimization problems with deep learning, in particular combinatorial optimization problems, where the optimal solutions must be found in a very large but finite space of possible configurations. The size of this space scales exponentially with the input size of the problem, thus making an exhaustive search-based solution strategy unfeasible. Well-known examples are solving the [Rubik's cube](#), the travelling salesman problem or finding the 3D structure of [proteins](#). An [Article](#) in this issue by Mohammed Hibat-Allah et al. proposes a new approach in optimization, fusing concepts from classical and quantum physics

and from deep learning, of particular use for finding the ground state of disordered Ising systems.

In statistical physics, optimization problems can be tackled with a computational approach called simulated annealing. This is a heuristic process inspired by annealing in metallurgy where a material is rapidly heated and subsequently slowly cooled so that the material can rearrange at the microscale into an optimal configuration with desired properties. A simulated annealing algorithm explores an energy landscape to find its global minimum by gradually decreasing 'thermal fluctuations' (see the figure for an example of a rough energy landscape with a clear global minimum). Simulated annealing, both in its [classical](#) and [quantum](#) formulation, is widely useful for optimization problems, but the process of 'cooling down' (decreasing the thermal fluctuations) to explore the optimization landscape is generally a slow process. Hibat-Allah et al. combine simulated annealing with a so-called variational approach, by parameterizing the joint distribution of the system's state via a recurrent neural network. The new method, dubbed variational neural annealing, significantly speeds up the simulated annealing process.

Many more fruitful interactions between physics and machine learning can be expected. There is much excitement around the promise of merging machine learning with quantum information approaches. One focus is [accelerating machine learning](#) with quantum computers — betting on the availability of reliable quantum hardware in the near future. Vice versa, concepts from machine learning are fuelling [advances](#) in quantum computing. Interested readers can find further inspiration in the comprehensive [Review article](#) from Giuseppe Carleo et al. or tune into the upcoming NeurIPS [workshop](#) on machine learning and the physical sciences. □

Published online: 16 November 2021  
<https://doi.org/10.1038/s42256-021-00416-w>