

# Learning without the machine: Thermodynamic, Chemical, Evolutionary

**Alexei Tkachenko**<sup>1</sup>

<sup>1</sup>Brookhaven National Laboratory, Upton, United States

Recent advances in artificial intelligence have been driven by increasingly large digital neural networks, but their remarkable performance comes with a growing cost in energy, hardware complexity, and centralized computation. This talk explores an alternative perspective: learning as a physical process that can emerge directly from matter itself, without relying on conventional digital machinery. I will discuss three complementary routes to such "learning without the machine," spanning thermodynamic, chemical, and evolutionary frameworks.

First, I will present a thermodynamic description of quasi-static analog deep neural networks [1]. In this approach, modern feedforward architectures can be mapped onto an effective free-energy functional, giving neural computation a direct statistical-mechanical interpretation. Inference corresponds to relaxation toward a unique free-energy minimum in which all network constraints are satisfied. Because this process can be carried out quasi-statically and reversibly, the minimal thermodynamic cost of inference can, in principle, vanish. This sharply contrasts with digital hardware, where bit erasure and metastability impose Landauer-type bounds on energy dissipation. Training, however, is fundamentally different. When both inputs and target outputs are simultaneously imposed, the system becomes overconstrained, generating mechanical-like stresses that propagate backward through the network. These stresses recover the structure of backpropagation, while slow annealing of weights and biases relaxes them, producing learning as a purely physical process rather than as explicit optimization of an externally imposed loss function. Within this framework, one obtains a universal lower bound on training energy that scales extensively with both the number of trainable parameters and the dataset size, revealing a fundamental asymmetry between nearly reversible inference and necessarily dissipative adaptation.

Second, I will turn to chemical learning in competitive dimerization networks, where molecules act as computational degrees of freedom [2]. In these systems, molecular species reversibly bind into pairwise complexes, and the collective equilibrium of the resulting network performs classification tasks. Inputs are encoded in molecular concentrations, outputs are read from the fugacities of designated reporter species, and binding affinities play the role of trainable couplings. Unlike standard neural networks, computation here is distributed, bidirectional, and embodied in the thermodynamics of reversible chemistry. This framework is especially attractive for DNA-based implementations, where sequence design provides direct control over interaction strengths. I will describe how such chemical networks can be trained through directed evolution rather than gradient descent. Mutation, amplification, selection, and recycling allow molecular populations to discover useful computational function directly in chemical space. Using multiclass classification of noisy input cocktails as a benchmark, these evolved networks achieve strong separation between target "on" and "off" outputs, high mutual information, and substantial robustness to noise. Remarkably, evolutionary search and conventional *in silico* optimization can reach comparable performance while identifying structurally diverse solutions, highlighting the richness and degeneracy of functional physical learning systems.

Together, these results suggest a broader paradigm in which learning can emerge from thermodynamic / chemical equilibrium, and evolutionary selection.

## References:

AV Tkachenko, "Thermodynamic bounds on energy use in quasi-static Deep Neural Networks", arXiv:2503.09980 (2025)

A.V. Tkachenko, B.M. Mognetti, S. Maslov, "Evolutionary chemical learning in dimerization networks", arXiv:2506.14006 (2025)

