CS 229br: Foundations of Deep Learning

Lecture 2: Generative Models

Boaz Barak

Gustaf Ahdritz

Gal Kaplun

Plan: Generative models

- Energy based models
- Sampling: Metropolis-Hastings and Langevin dynamics
- Digression: statistical physics, variational principle, free energy
- Encoder-Decoder Architecture
- PCA
- Variational Auto Encoder
- Normalizing Flow

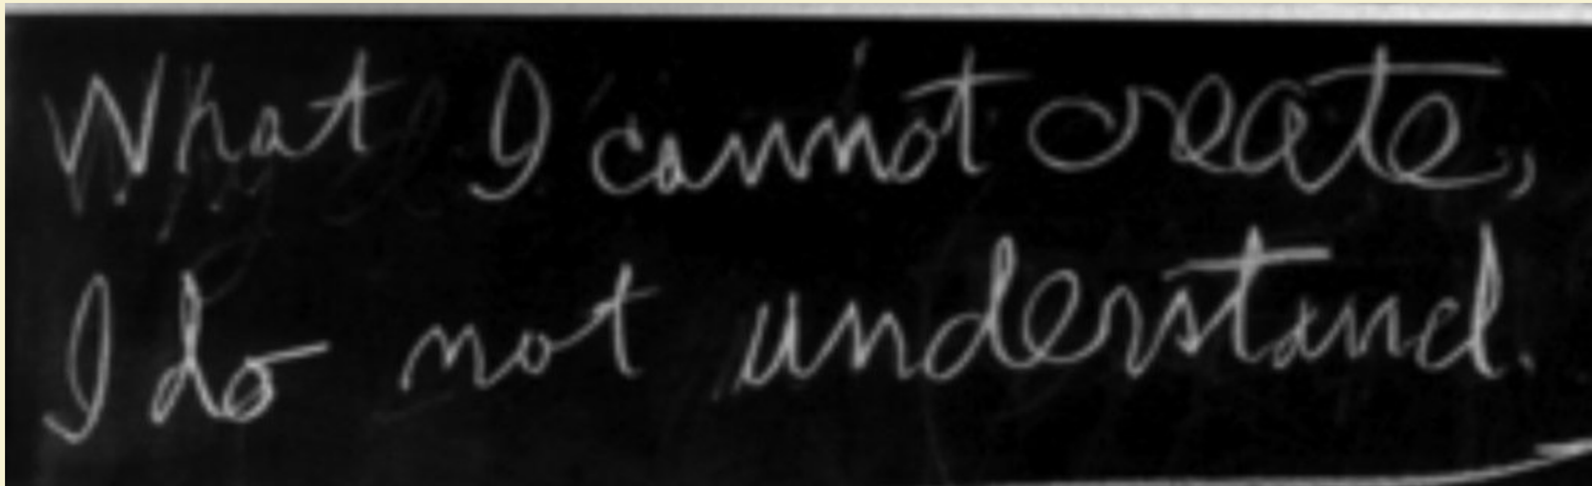

Energy based models

Proc. Natl. Acad. Sci. USA
Vol. 79, pp. 2554–2558, April 1982
Biophysics

Neural networks and physical systems with emergent collective computational abilities

(associative memory/parallel processing/categorization/content-addressable memory/fail-soft devices)

J. J. Hopfield

Division of Chemistry and Biology, California Institute of Technology, Pasadena, California 91125; and Bell Laboratories, Murray Hill, New Jersey 07974

Contributed by John J. Hopfield, January 15, 1982

Reducing the Dimensionality of Data with Neural Networks

G. E. Hinton* and R. R. Salakhutdinov

High-dimensional data can be converted to low-dimensional codes by training a multilayer neural network with a small central layer to reconstruct high-dimensional input vectors. Gradient descent can be used for fine-tuning the weights in such "autoencoder" networks, but this works well only if the initial weights are close to a good solution. We describe an effective way of initializing the weights that allows deep autoencoder networks to learn low-dimensional codes that work much better than principal components analysis as a tool to reduce the dimensionality of data.

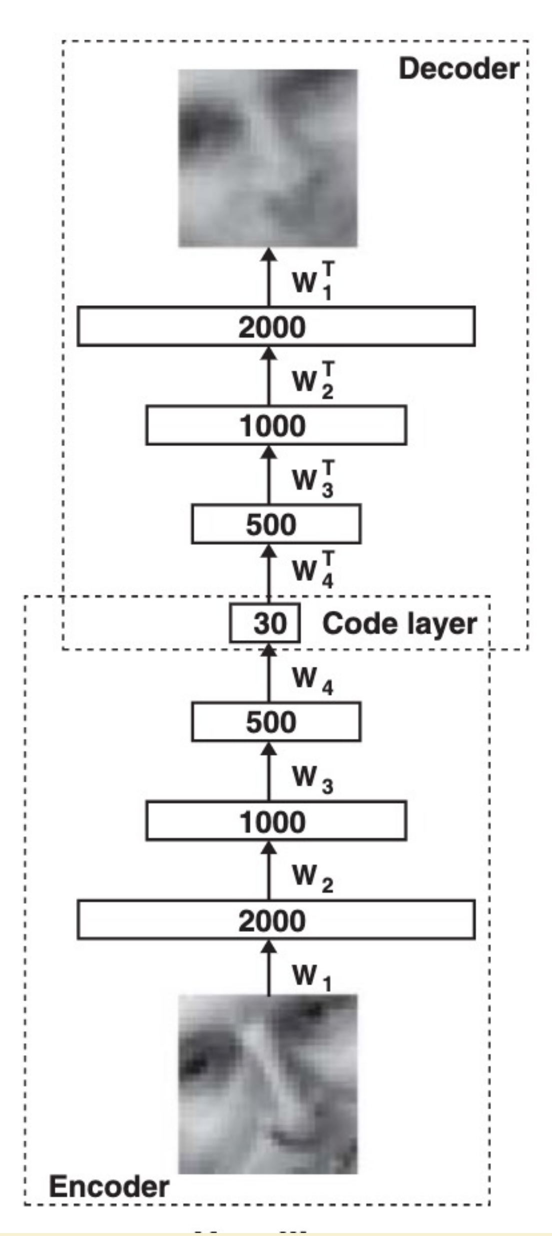

Energy based models

Implicit Generation and Modeling with Energy-Based Models

Yilun Du *
MIT CSAIL

Igor Mordatch
Google Brain

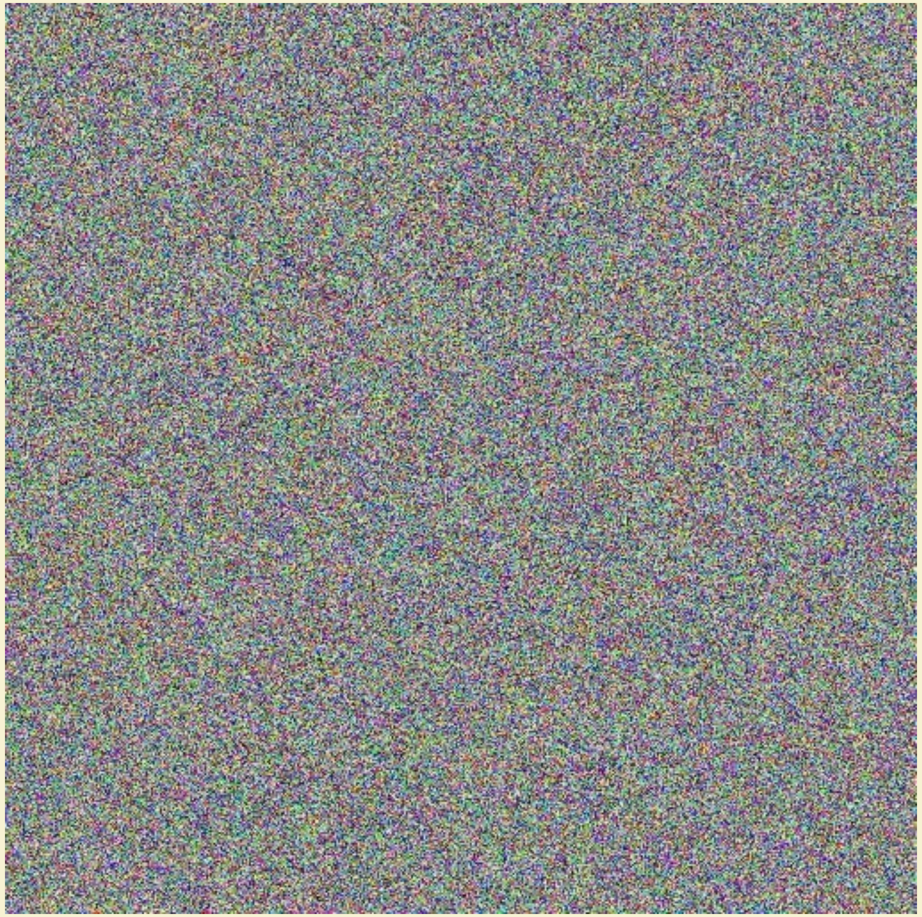

Energy based models

Meta's AI luminary LeCun explores deep learning's energy frontier

So-called energy-based models, which borrow concepts from statistical physics, may lead the way to 'abstract prediction,' says Yann LeCun, allowing for a 'unified world model' for AI capable of planning.

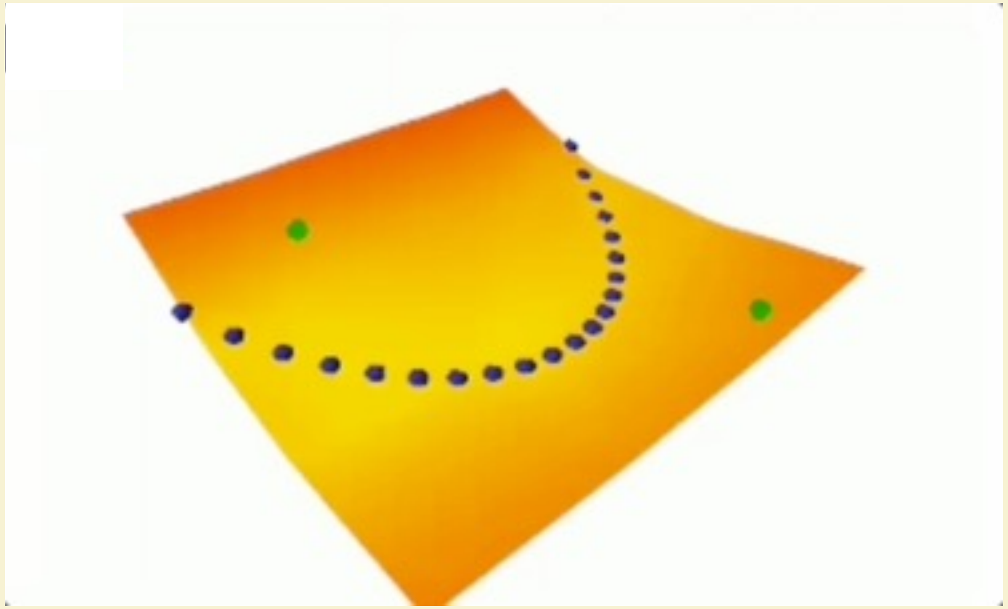

Move to whiteboard – see notes

Langevin dynamics

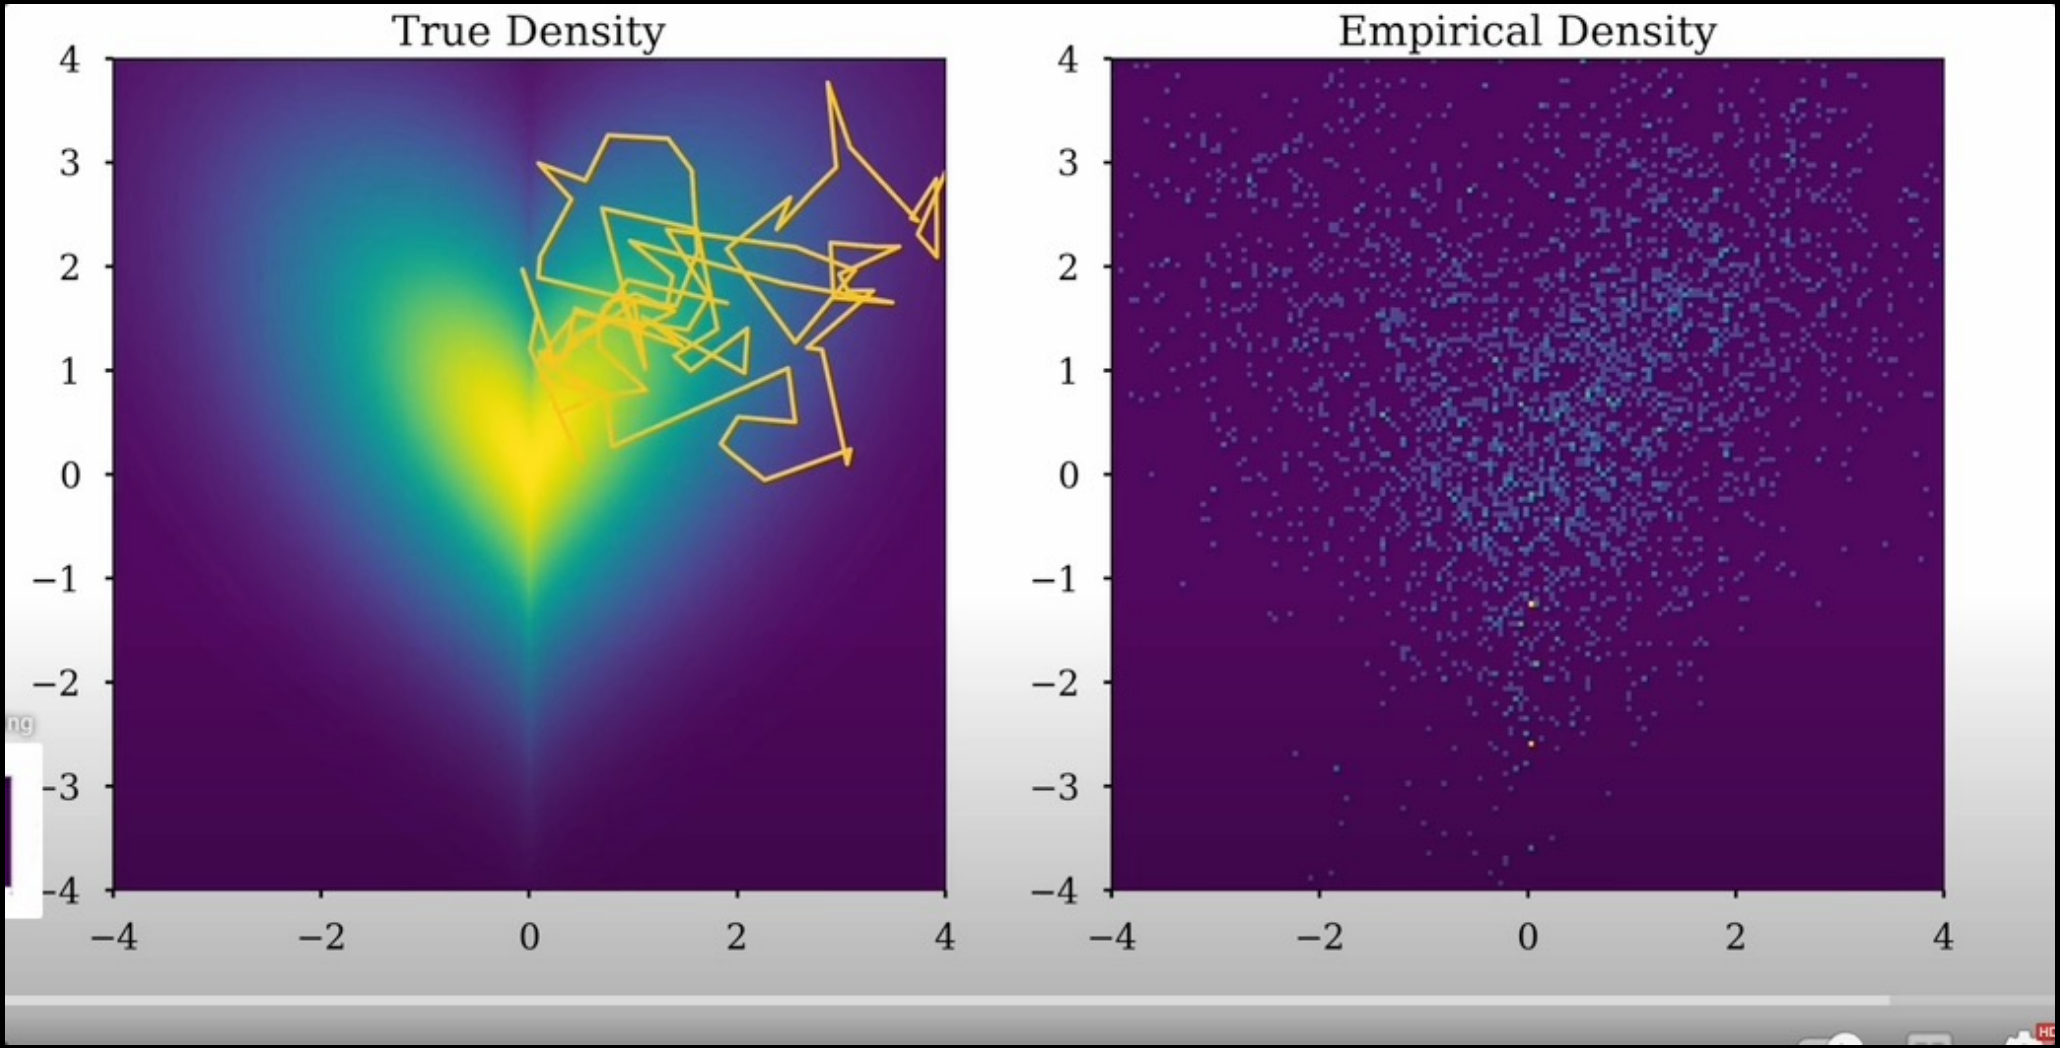

Abdul Fatir Ansari

Move to whiteboard – see notes

Water Ice

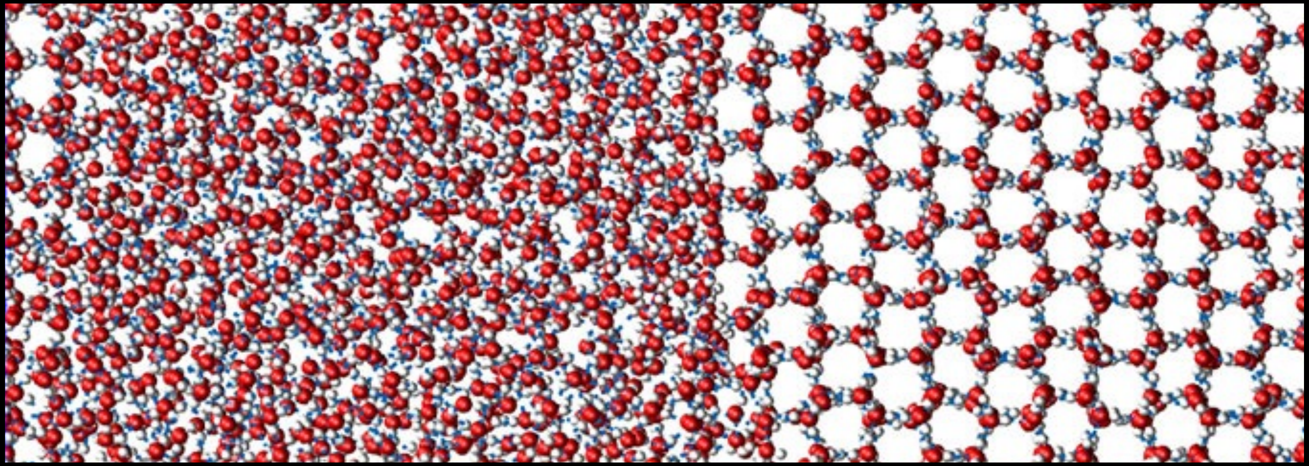

Lower energy

Higher temperature / entropy

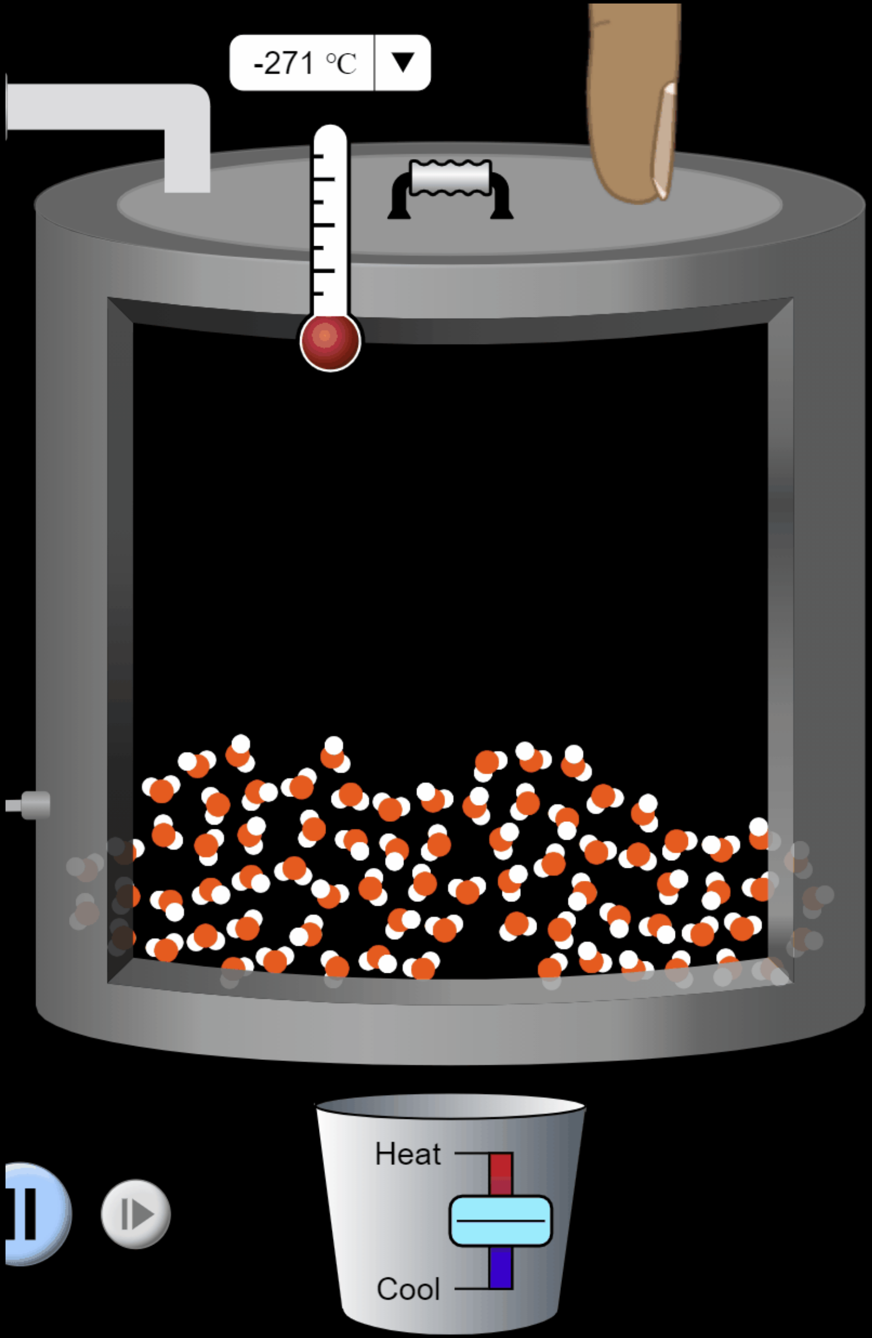

Ising model

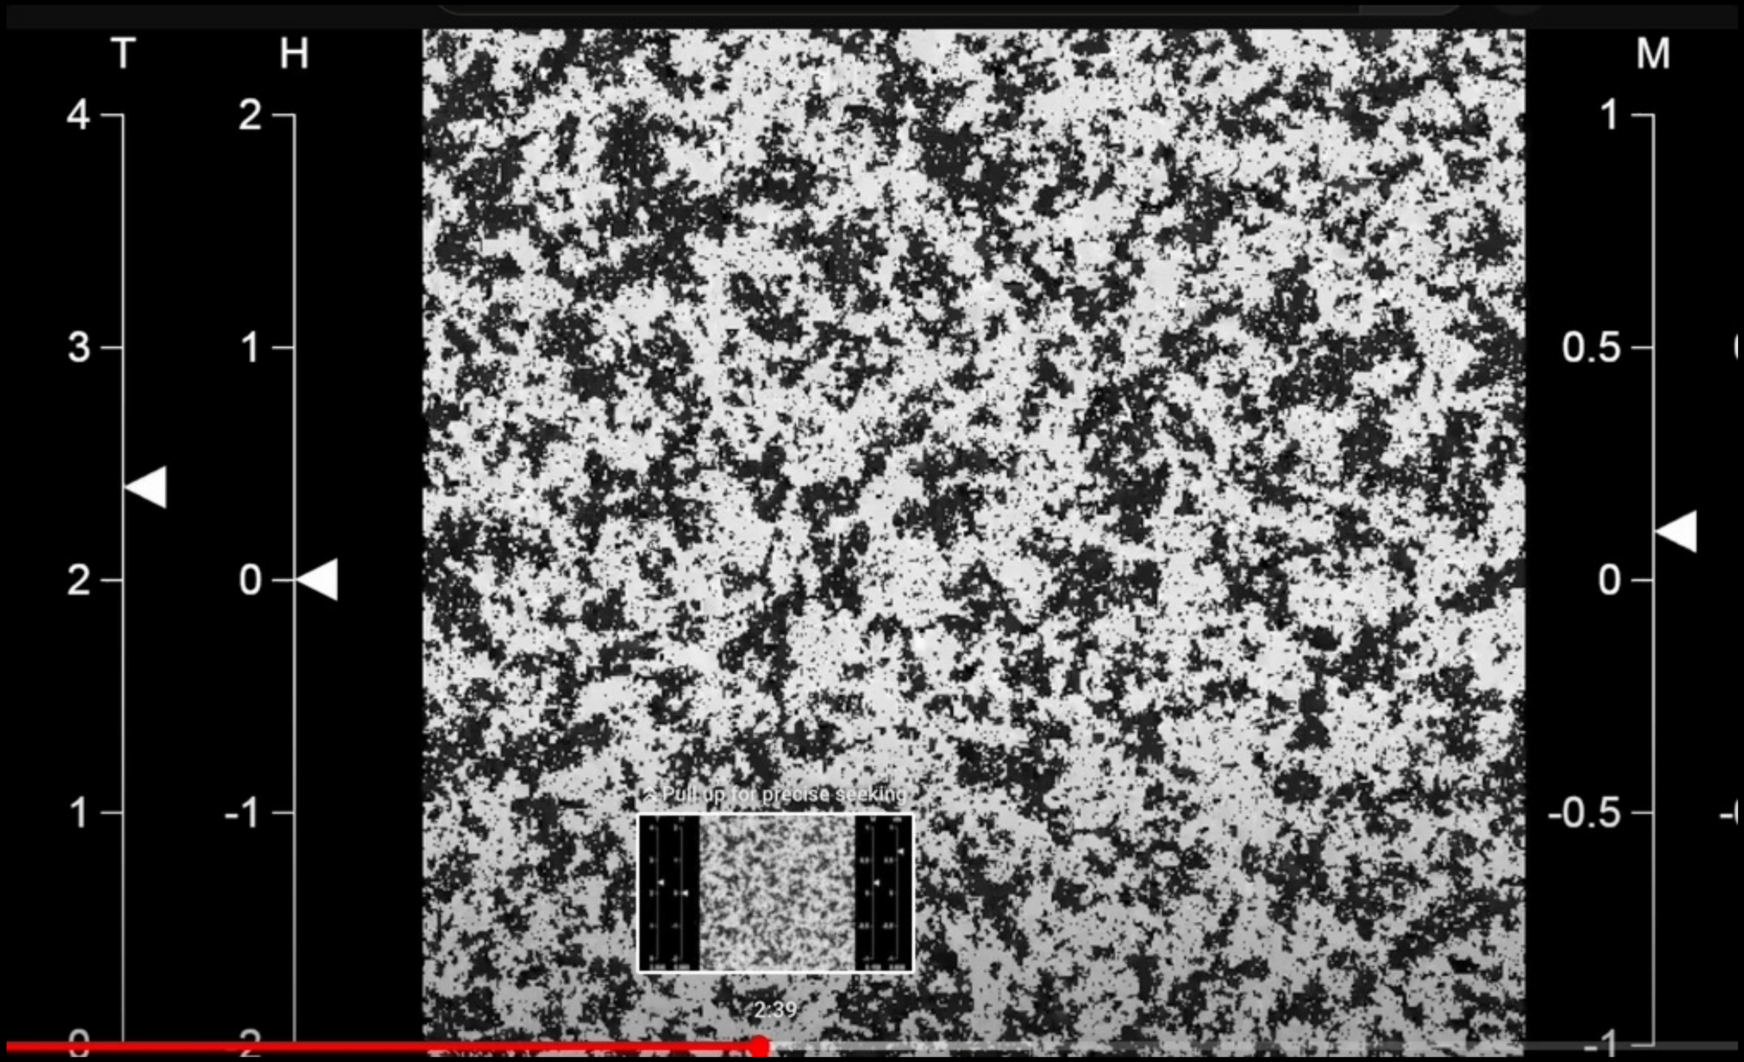

Principle of least action

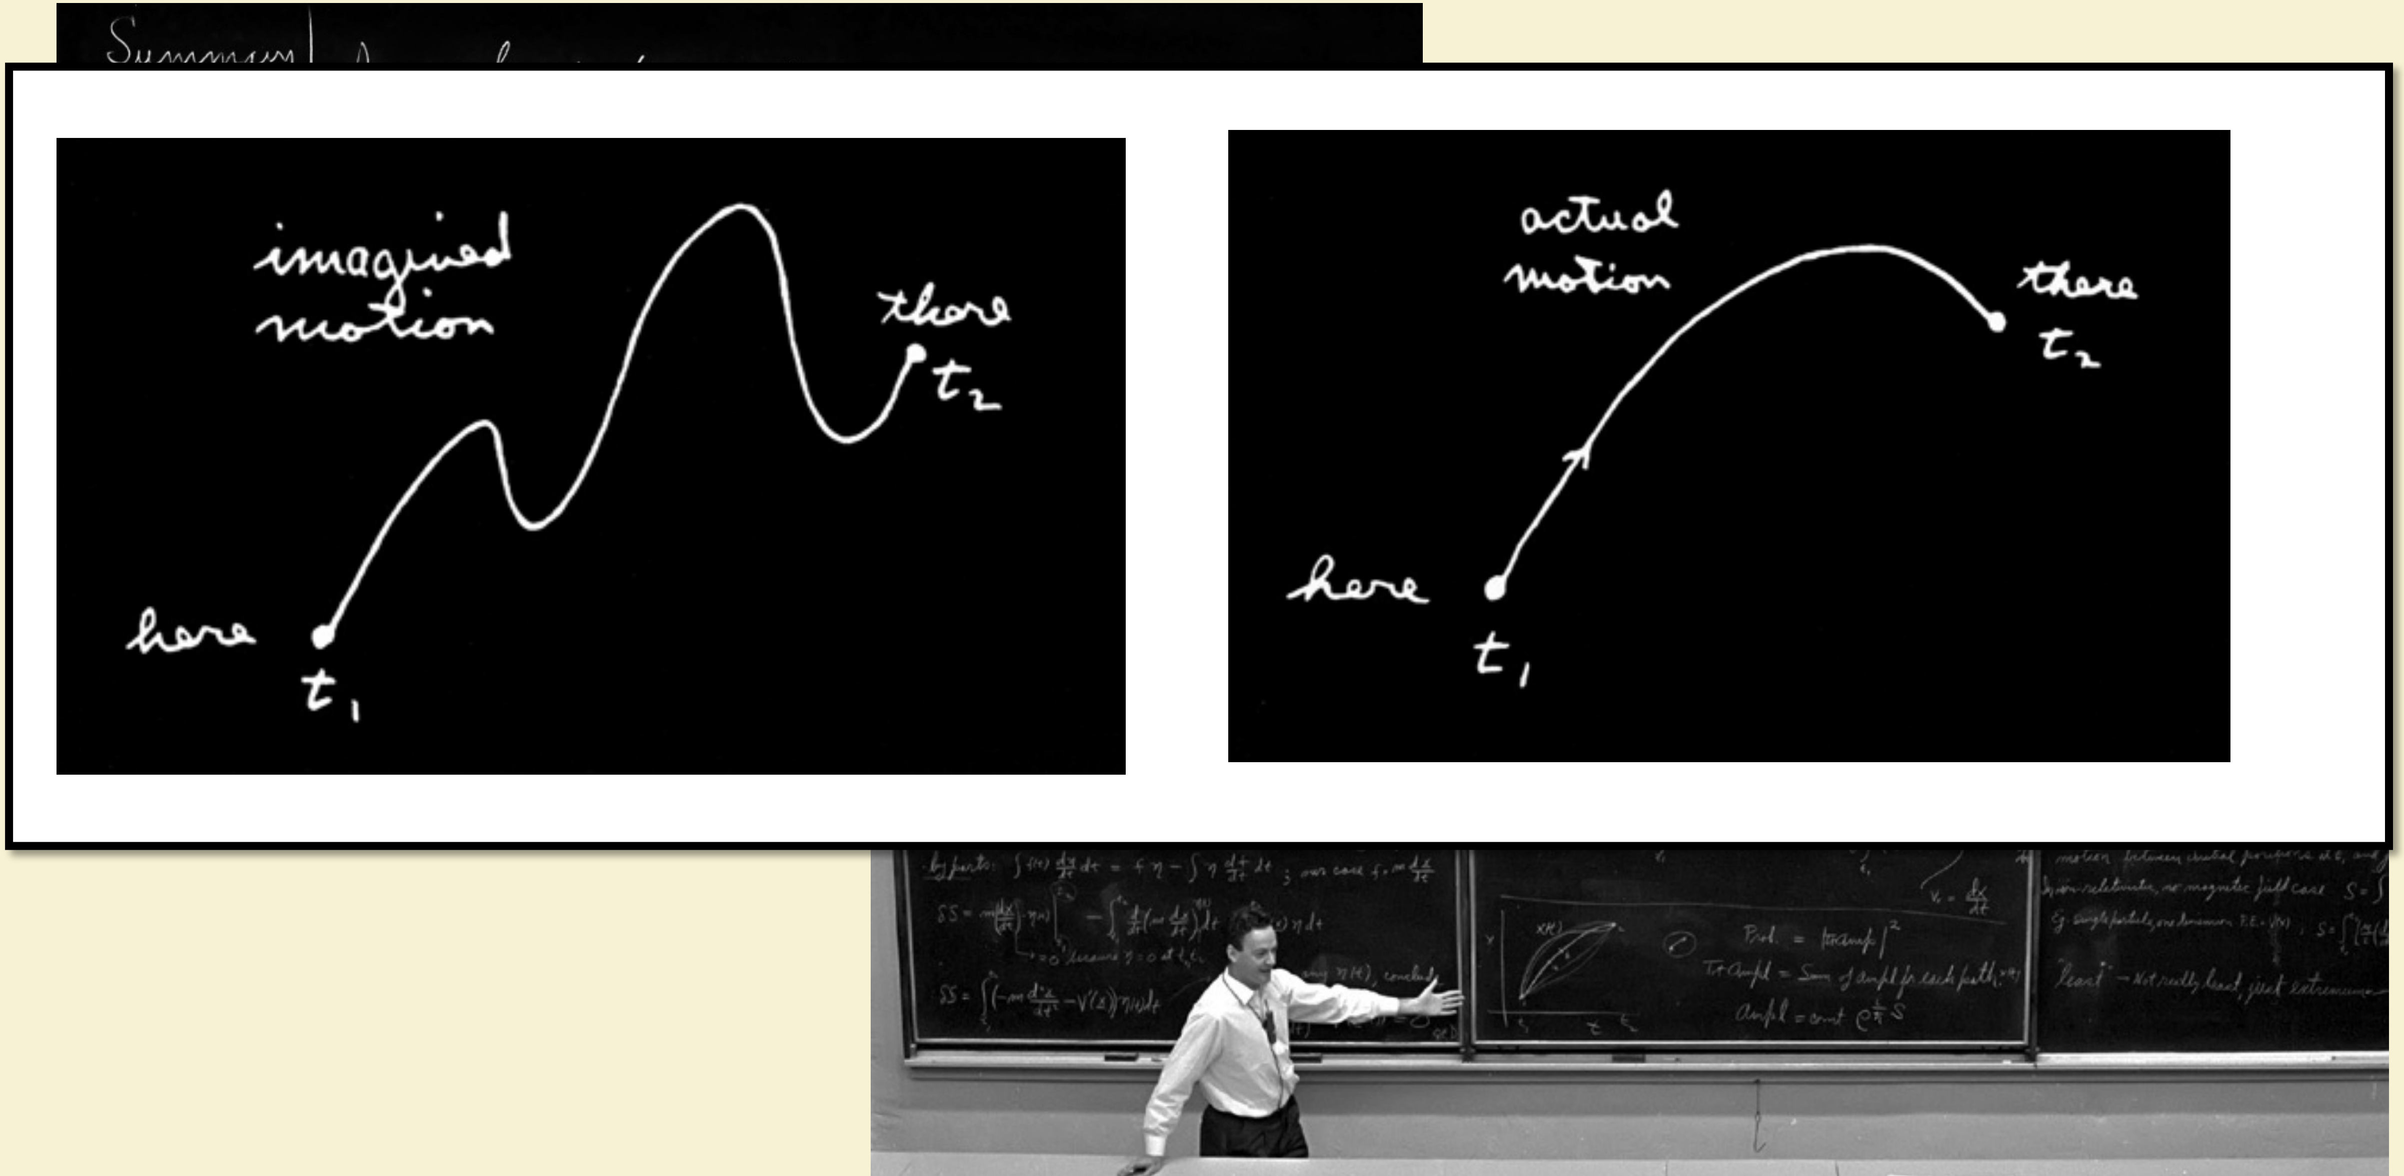

Two views of physical systems

"Local": State progresses from time t to $t+1$ based on local forces

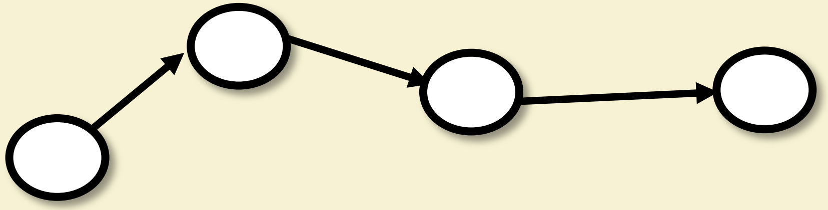

$$f: \mathbb{R}^d \to \mathbb{R}^d$$

$$s_{t+1} = s_t + f(s_t)$$

$$\frac{ds}{dt} = \mathrm{f}$$

"Global": System minimizes a global objective

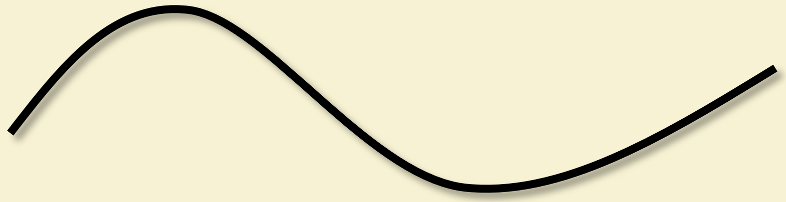

$\mathcal{L}$: Potential energy

$$f = -\nabla \mathcal{L}$$

Principle of least action:

s minimizes $\sum_t |\mathcal{L}(s_{t+1}) - \mathcal{L}(s_t)|^2$