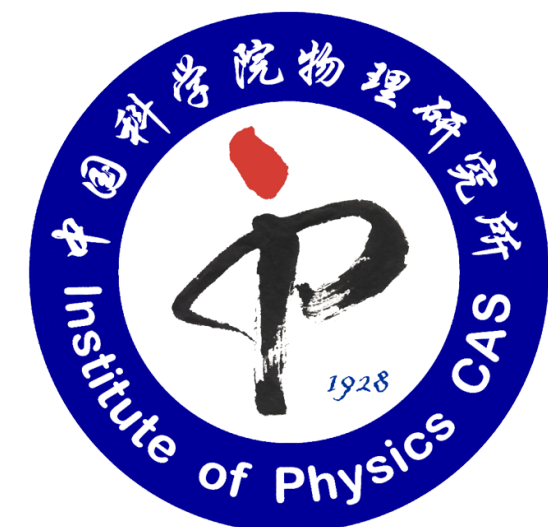


Generative AI for Science

Lei Wang (王磊)

Institute of Physics, CAS

<https://wangleiphy.github.io>



Discriminative learning



$$y = f(\mathbf{x})$$

or $p(y | \mathbf{x})$

Generative learning



$$p(\mathbf{x}, y)$$

Generative AI: a new buzz word in silicon valley

A Coming-Out Party for Generative A.I., Silicon Valley's New Craze

A celebration for Stability AI, the start-up behind the controversial Stable Diffusion image generator, represents the arrival of a new A.I. boom.

New York Times

Kevin Roose

Oct. 21, 2022

Protocol

Biz Carson

October 21, 2022

Sequoia's Sonya Huang: The generative AI hype is 'absolutely justified'

She's bullish on generative AI given the "superpowers" it gives humans who work with it.

<https://www.sequoiacap.com/article/generative-ai-a-creative-new-world/>

by Sonya Huang, Pat Grady and GPT-3

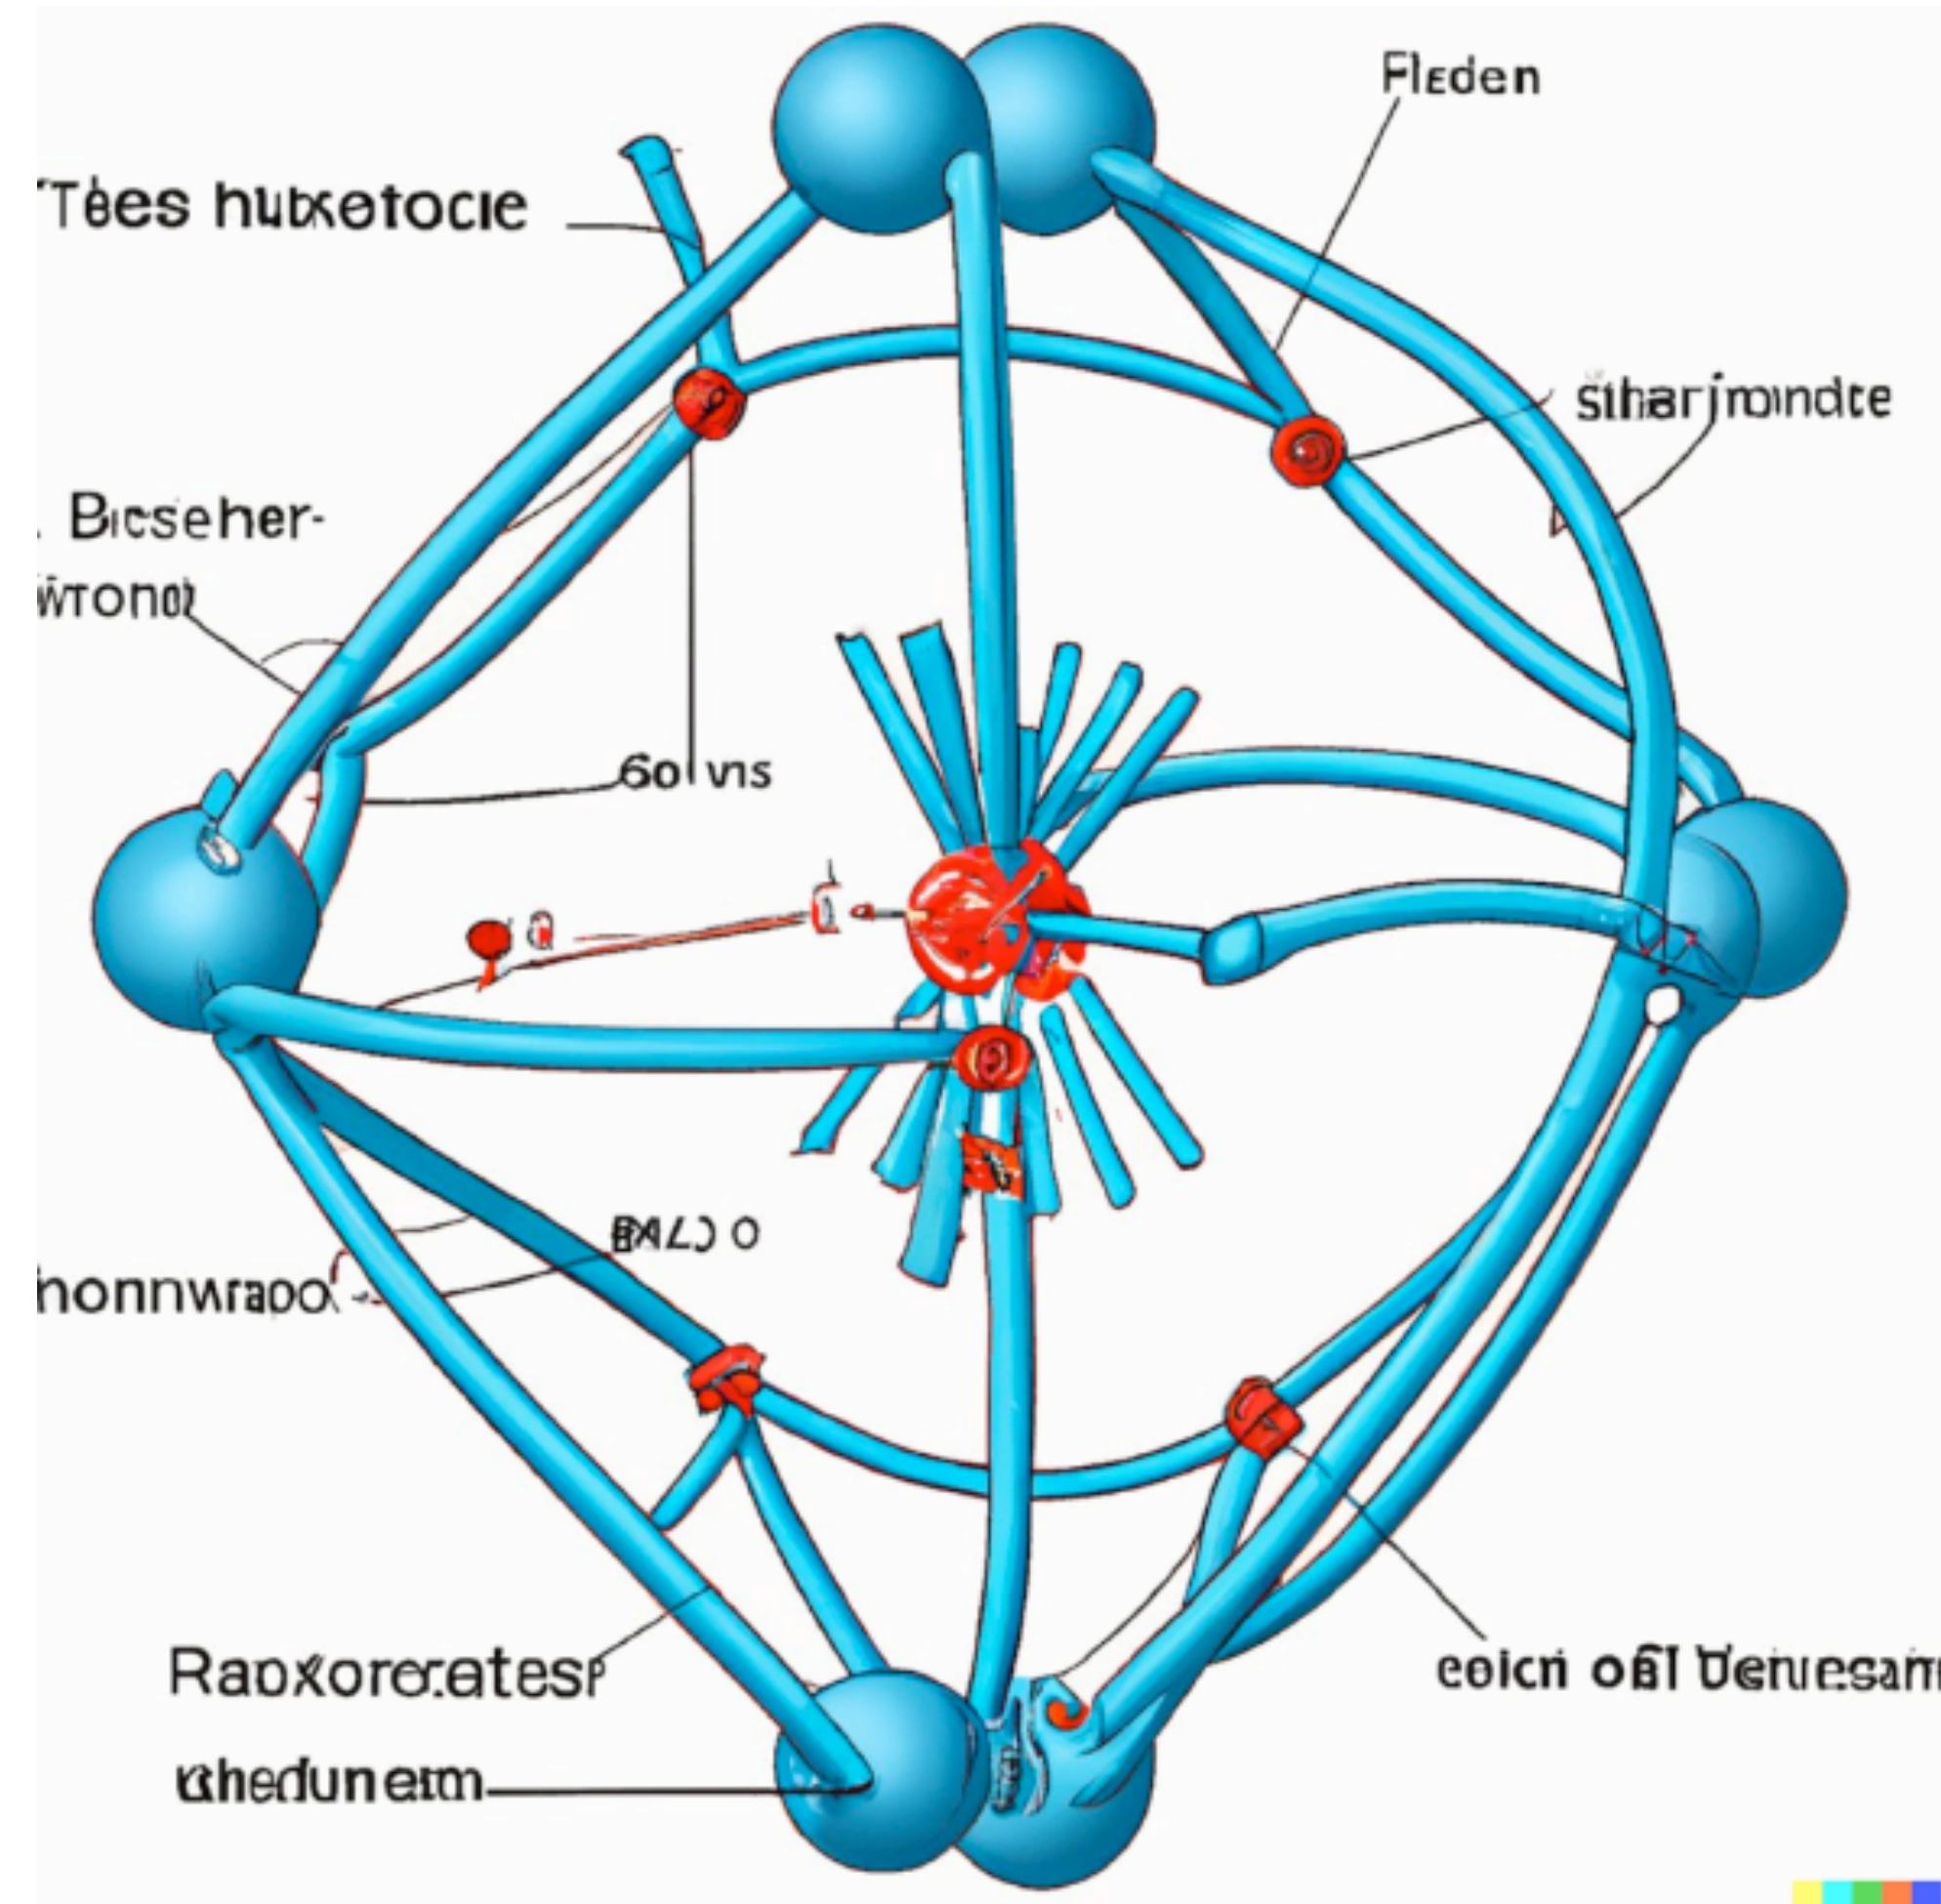
	PRE - 2020	2020	2022	2023?	2025?	2030?
TEXT	Spam detection Translation Basic Q&A	Basic copy writing First drafts	Longer form Second drafts	Vertical fine tuning gets good (scientific papers, etc)	Final drafts better than the human average	Final drafts better than professional writers
CODE	1-line auto-complete	Multi-line generation	Longer form Better accuracy	More languages More verticals	Text to product (draft)	Text to product (final), better than full-time developers
IMAGES			Art Logos Photography	Mock-ups (product design, architecture, etc.)	Final drafts (product design, architecture, etc.)	Final drafts better than professional artists, designers, photographers)
VIDEO / 3D / GAMING			First attempts at 3D/video models	Basic / first draft videos and 3D files	Second drafts	AI Roblox Video games and movies are personalized dreams

Large model availability: ● First attempts ● Almost there ● Ready for prime time

<https://huggingface.co/spaces/stabilityai/stable-diffusion>

the inner structure of an electron

Generate image



<https://future.com/how-to-build-gpt-3-for-science/>

How to Build a GPT-3 for Science *(scientific literature and data)*

Josh Nicholson

Posted August 18, 2022

Some examples of complex potential prompts are:

“Tell me why this hypothesis is wrong”

“Tell me why my treatment idea won’t work”

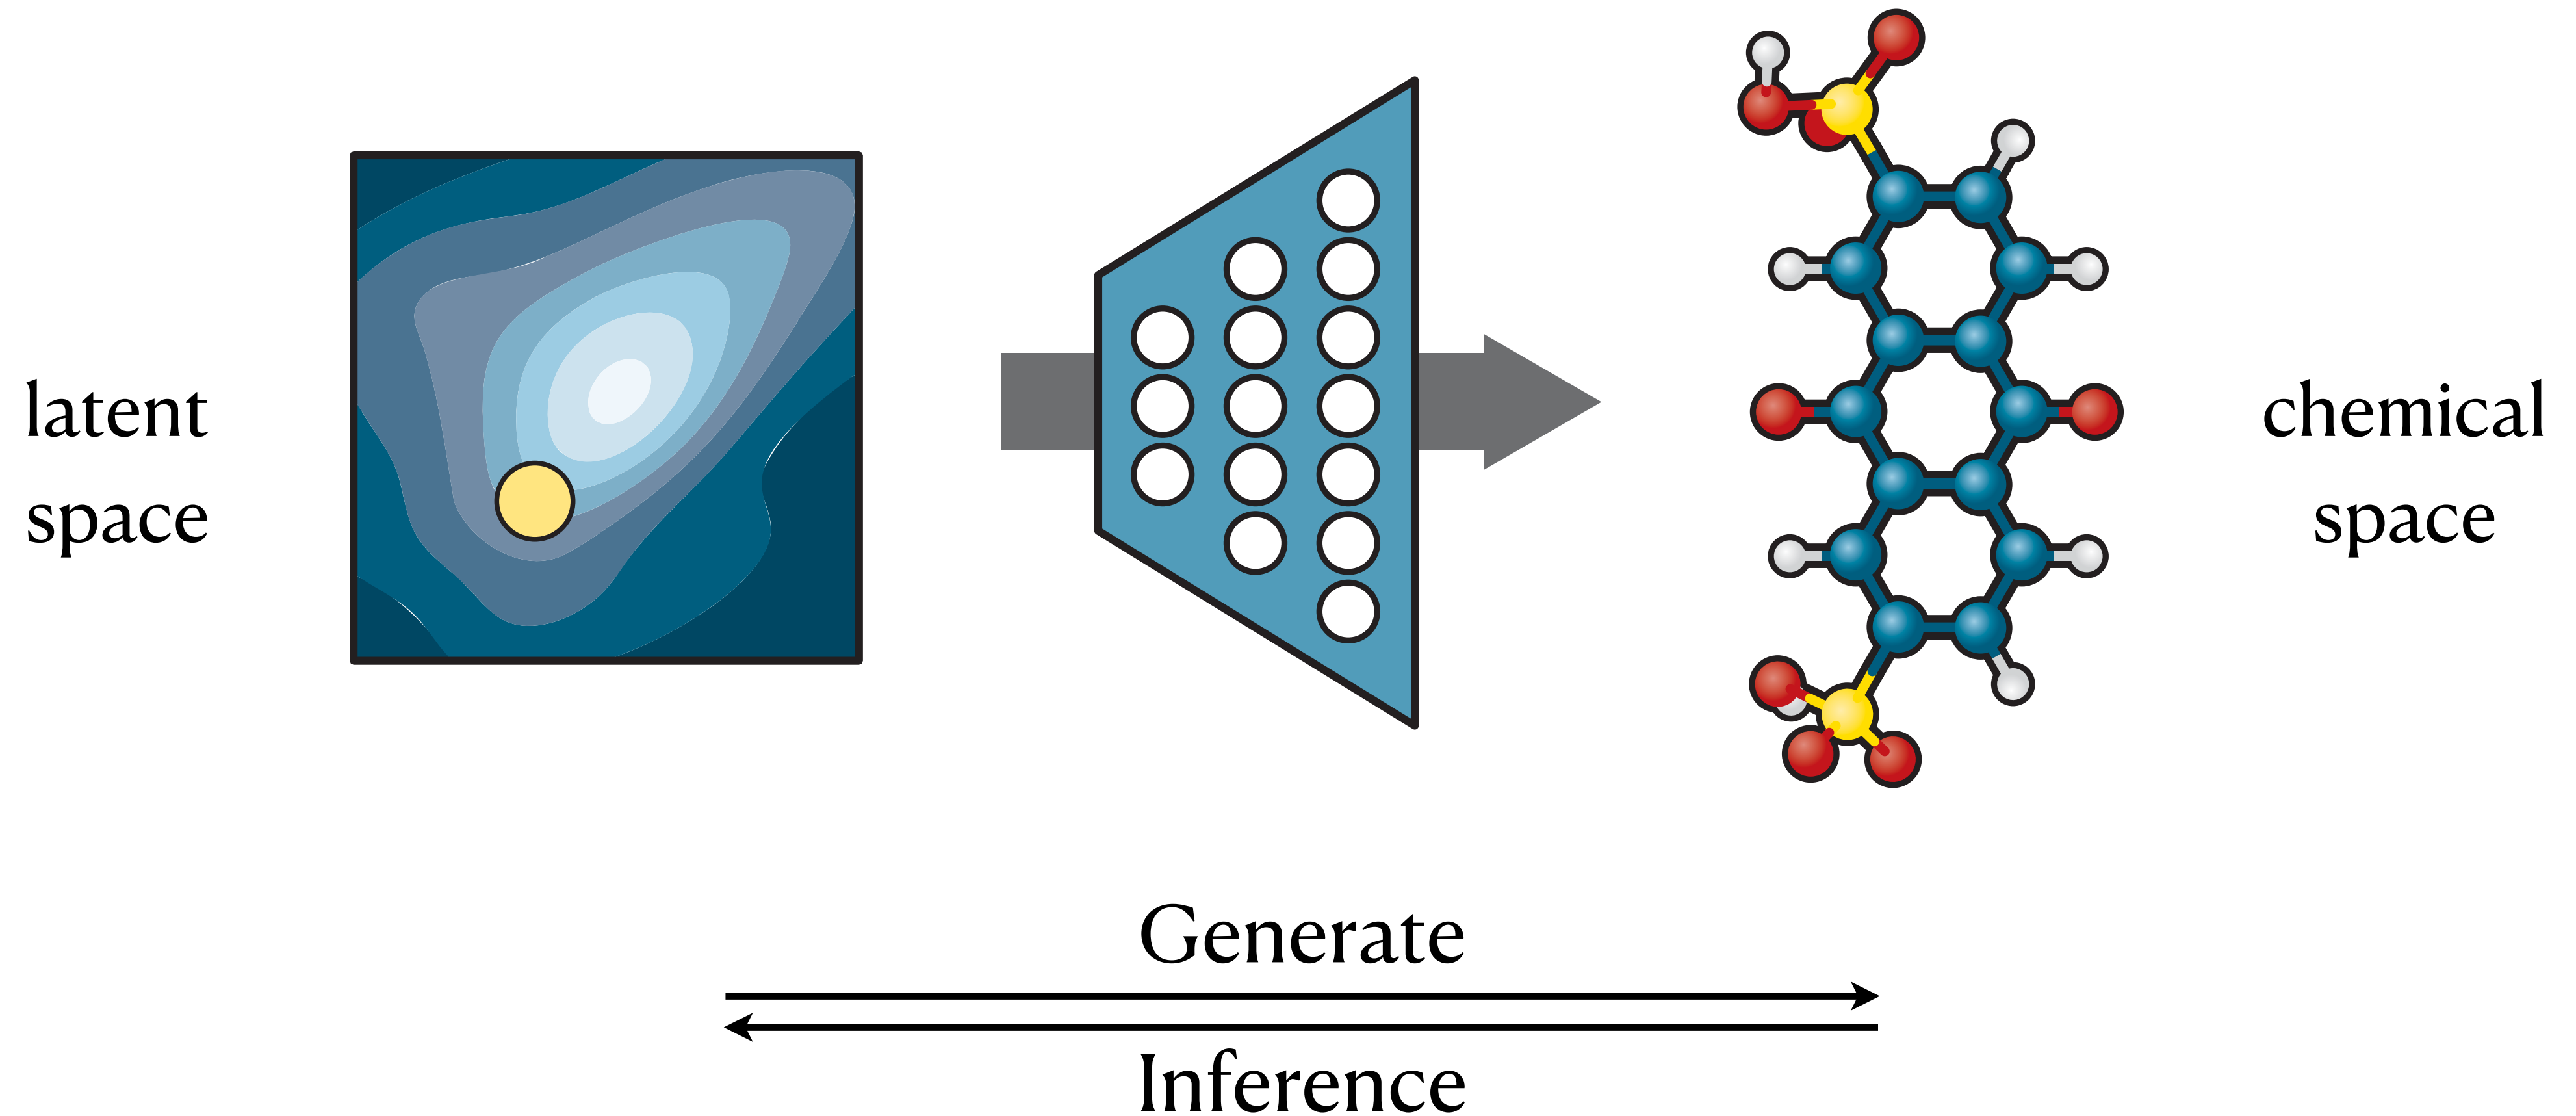
“Generate a new treatment idea”

“What evidence is there to support social policy X?”

“Who has published the most reliable research in this field?”

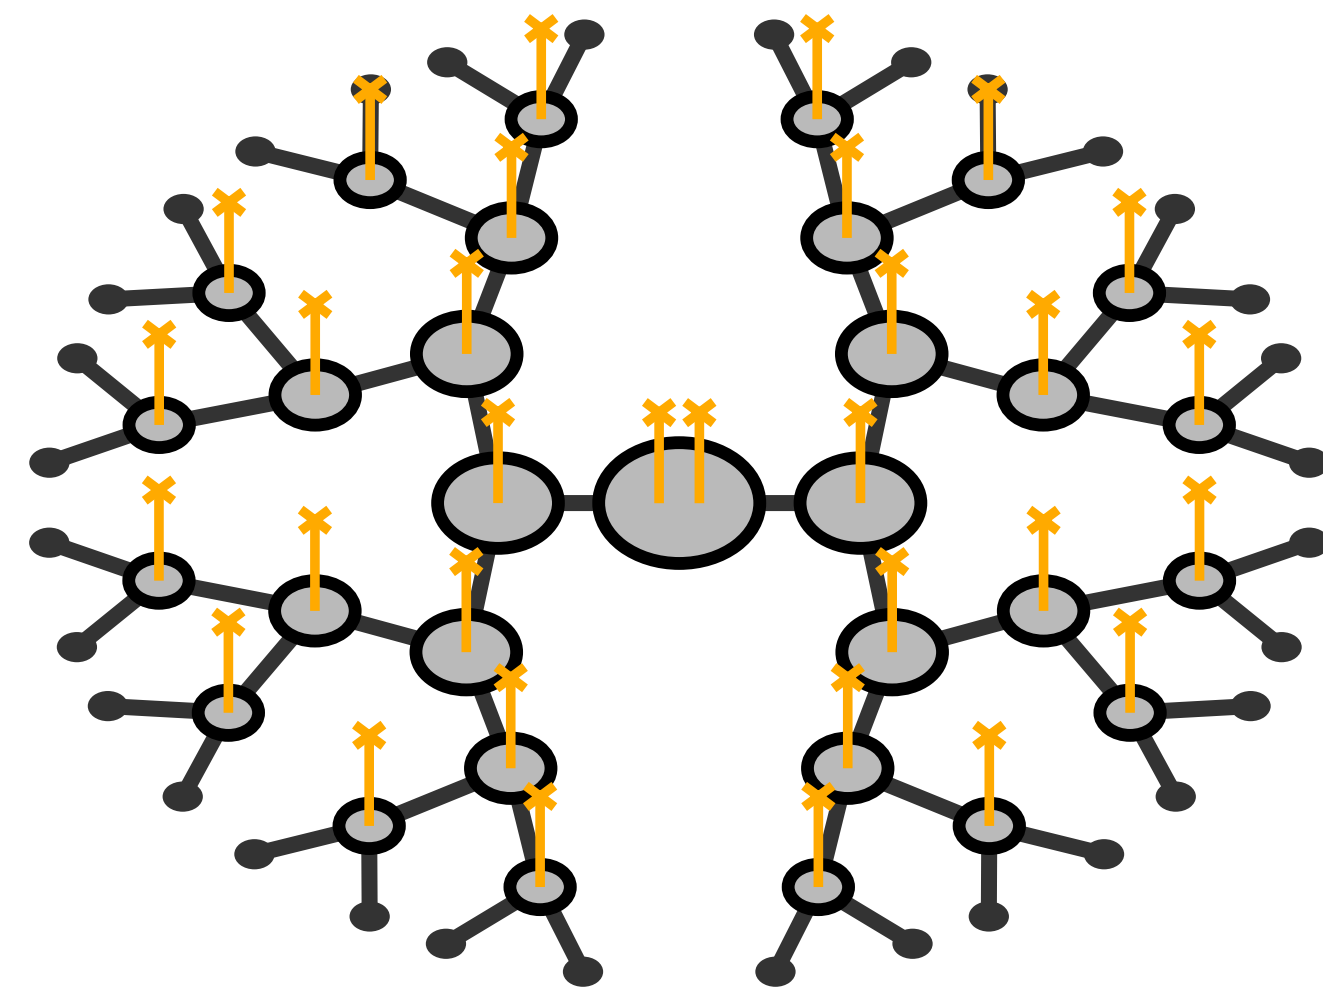
“Write me a scientific paper based on my data”

Generative AI for matter engineering



Generative AI for statistical physics

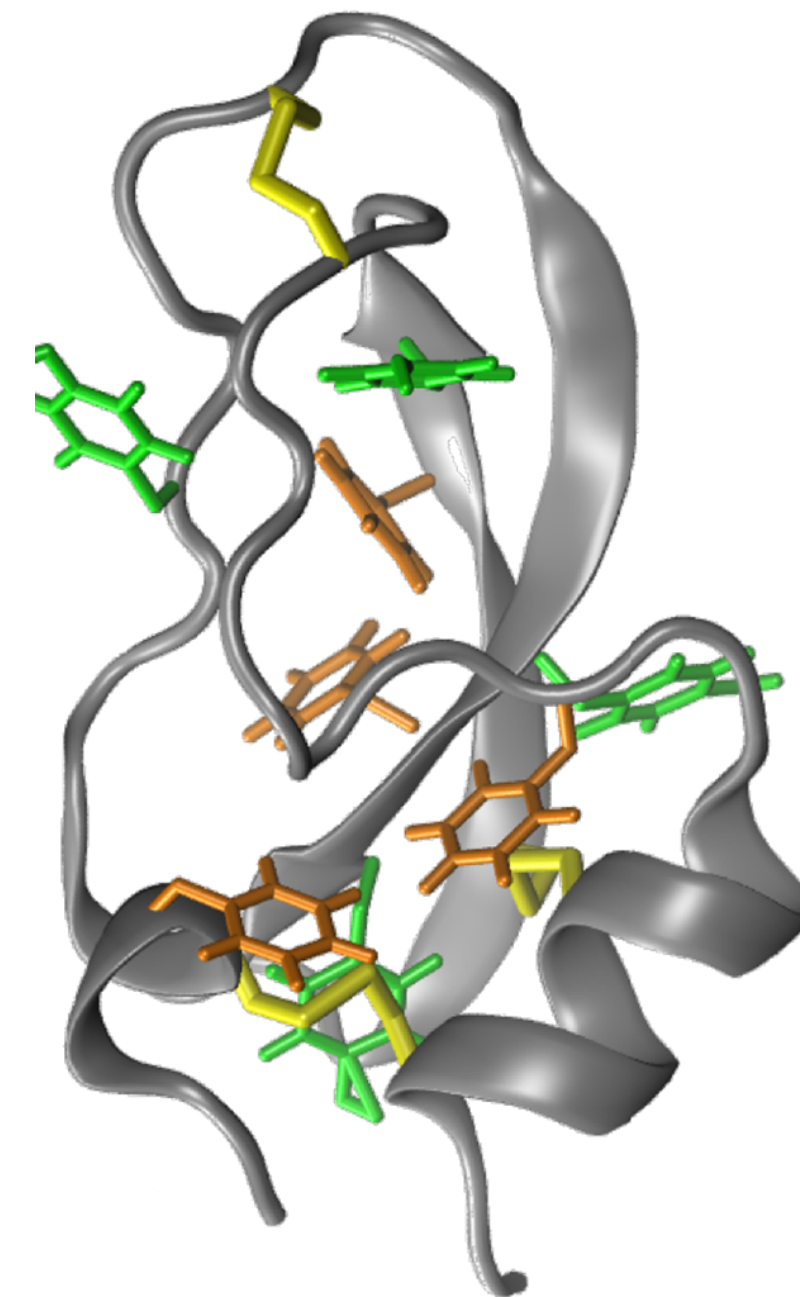
Renormalization group



Li and LW, PRL '18

Li, Dong, Zhang, LW, PRX '20

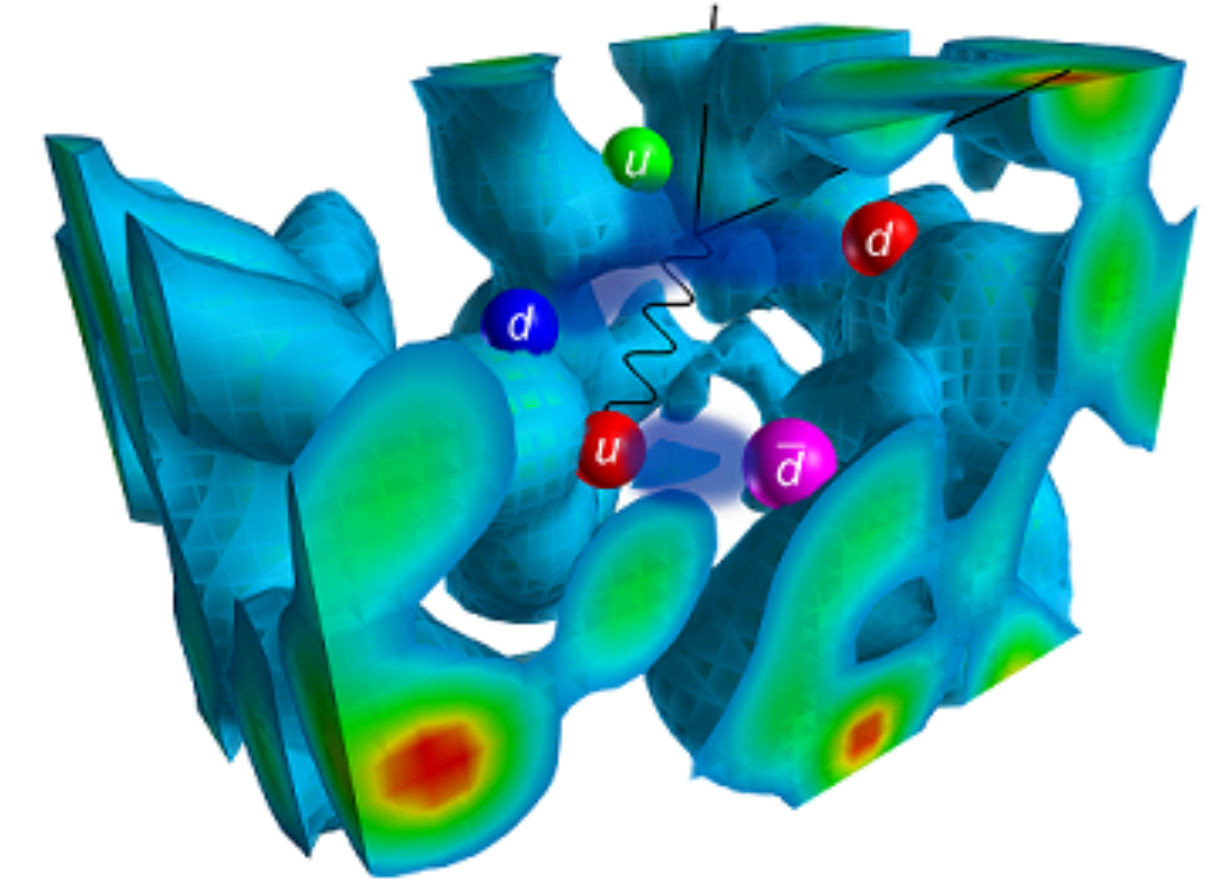
Molecular simulation



Noe et al, Science '19

Wirnsberger et al, JCP '20

Lattice field theory



Albergo et al, PRD '19

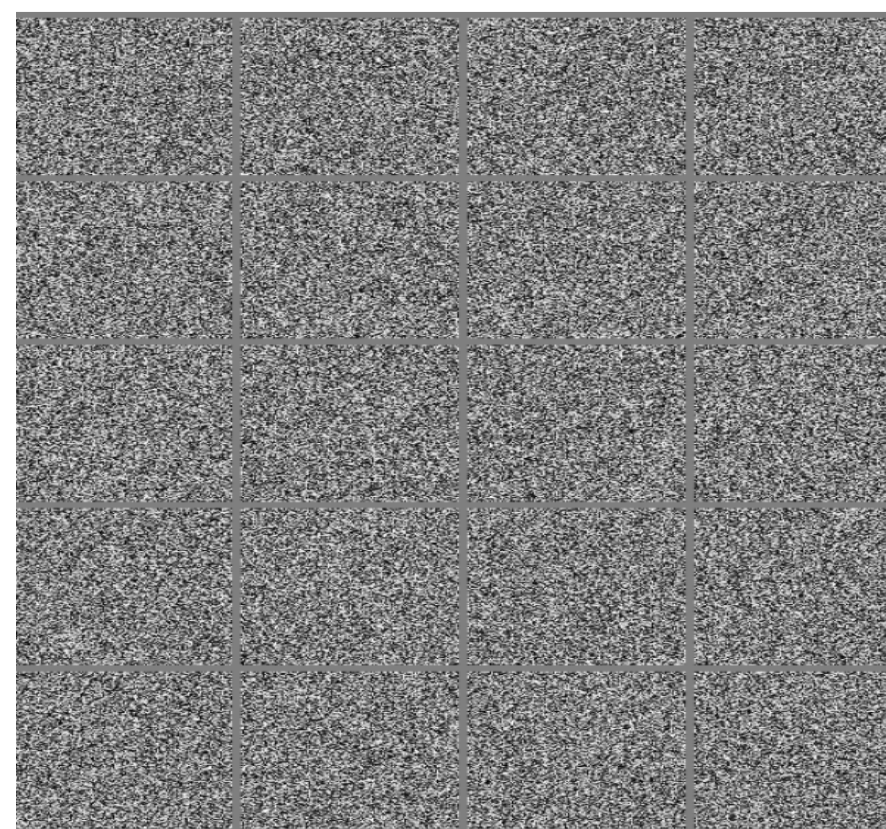
Kanwar et al, PRL '20

These are principled computation: quantitatively accurate, interpretable, reliable, and generalizable even without data

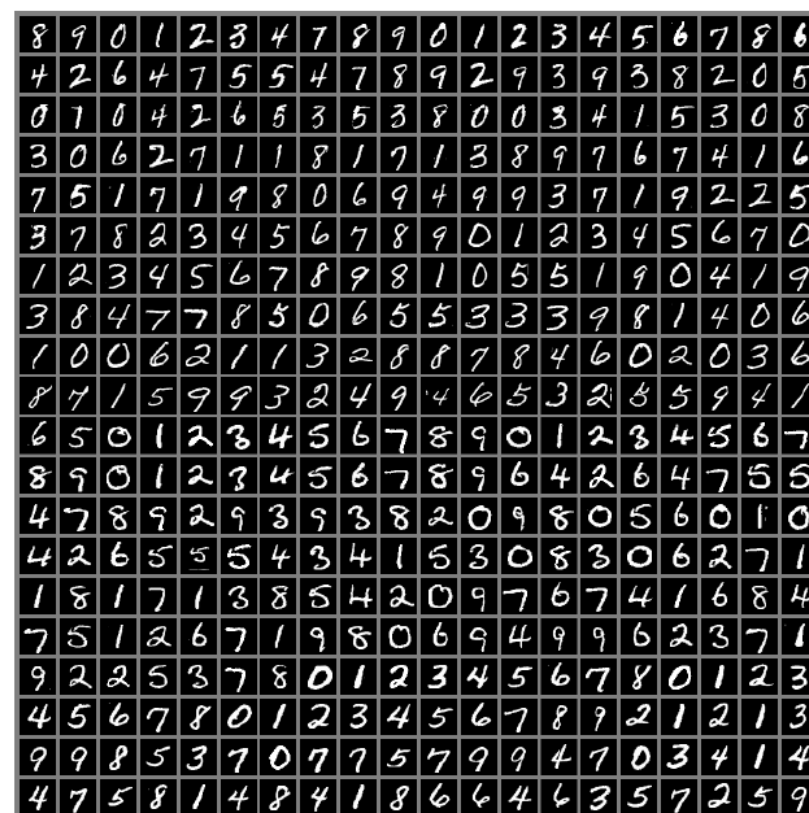
Probabilistic Generative Modeling

$$p(\mathbf{x})$$

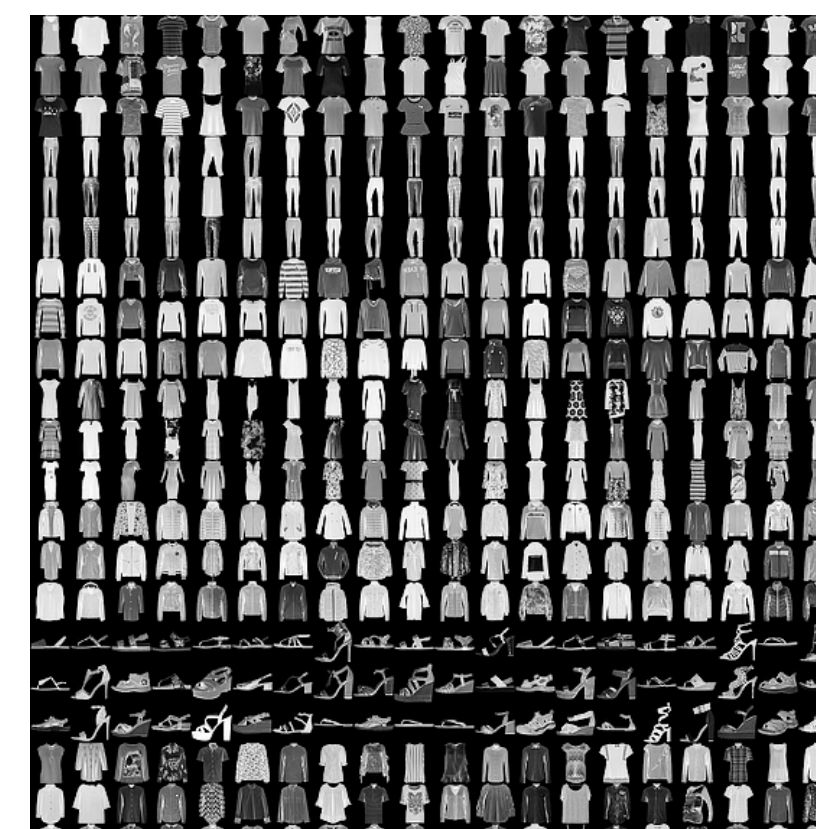
How to **express, learn, and sample from** a high-dimensional probability distribution?



“random” images



“natural” images



Probab

deling

DEEP LEARNING

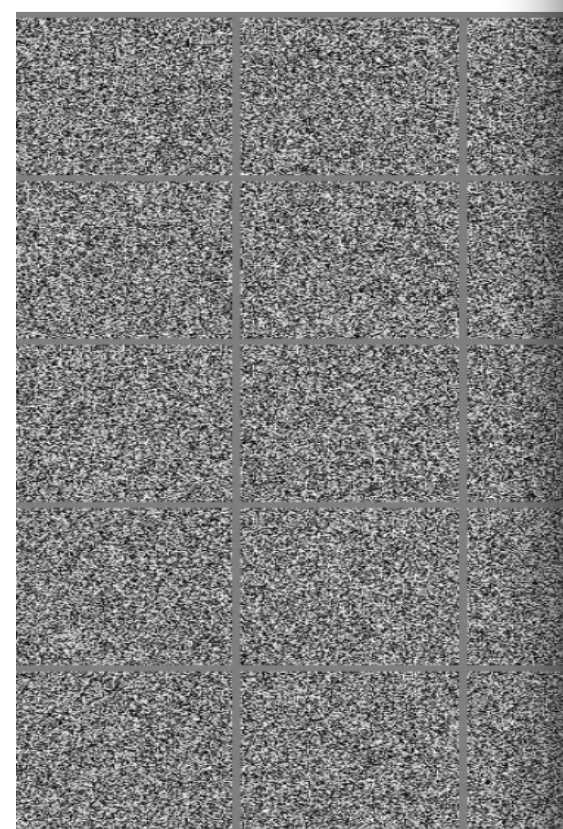
Ian Goodfellow, Yoshua Bengio,
and Aaron Courville

How to
high-dim

om a
ition ?

Page 159

*“... the images encountered in
AI applications occupy a
negligible proportion of
the volume of image space.”*



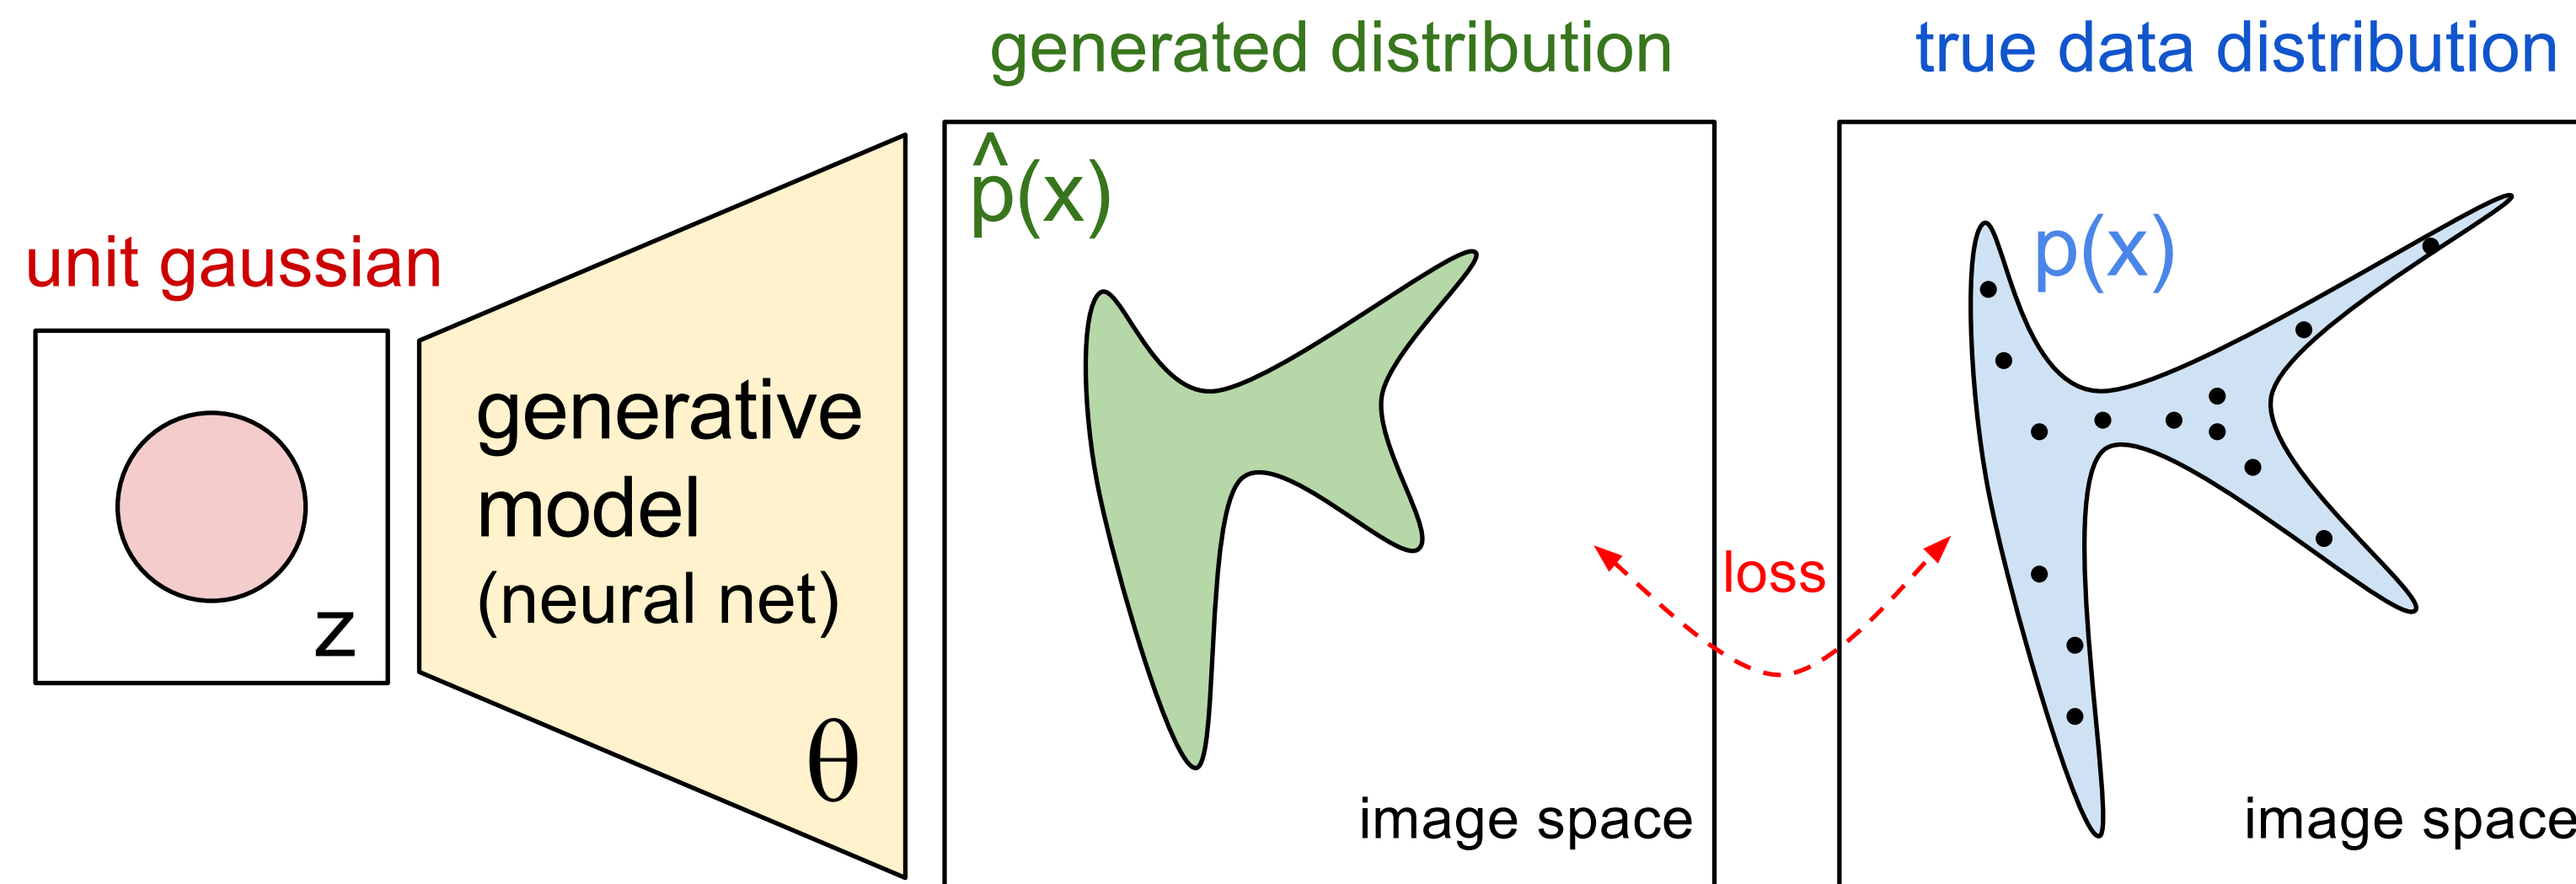
“random”



Probabilistic Generative Modeling

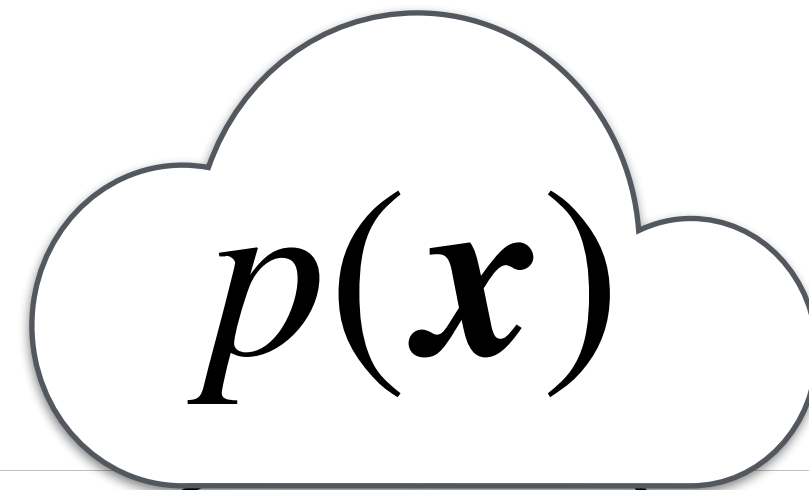
$$p(\mathbf{x})$$

How to **express, learn, and sample** from a high-dimensional probability distribution?



Generative models and their physics genes

Goodfellow,
NIPS tutorial, 1701.00160



Explicit density

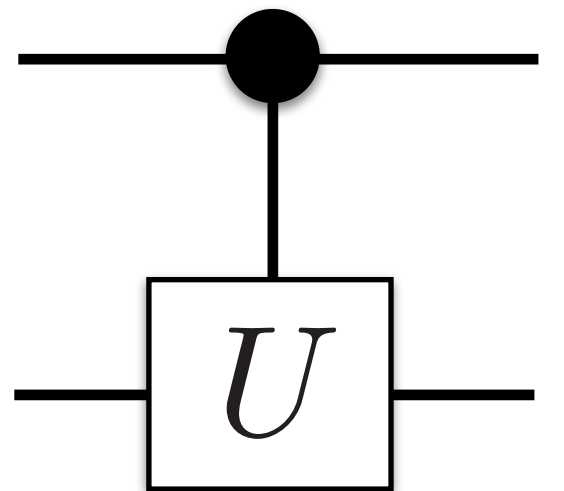
Implicit density

Direct
GAN

Tractable density

Approximate density

Markov Chain
GSN



Quantum
Circuits

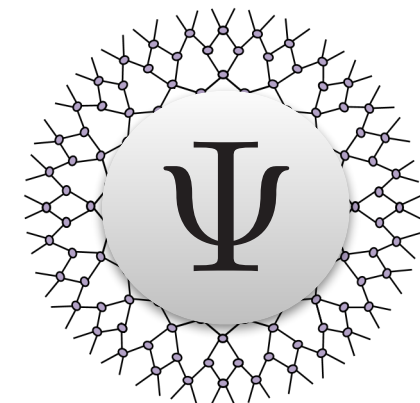
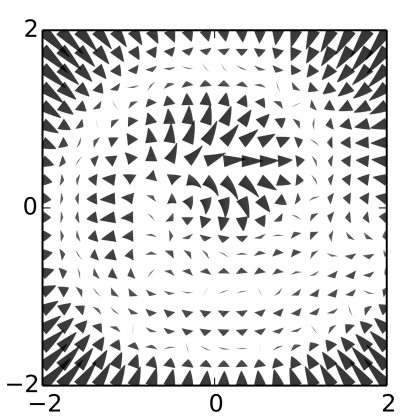
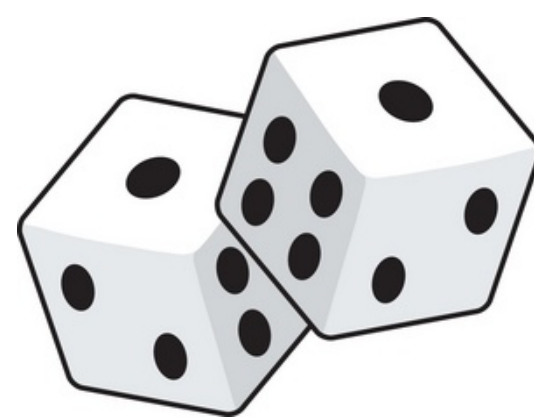
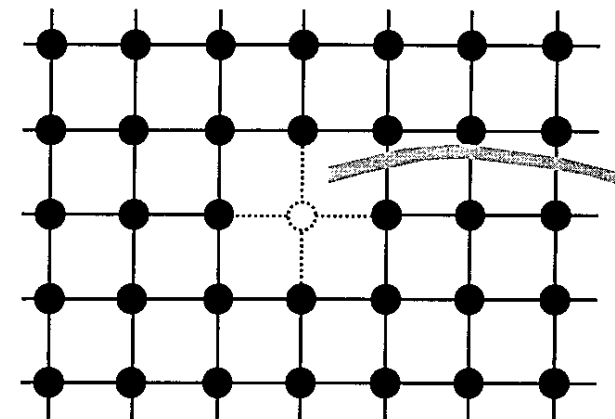
Liu, LW, PRA '18

- Fully visible belief nets
- NADE
- MADE
- PixelRNN
- Change of variables models (nonlinear ICA)

Variational

Markov Chain

Variational autoencoder Boltzmann machine + Diffusion models



Tensor
Networks

Han, Wang,
Fan, LW, Zhang
PRX '18

Generative modeling

Statistical physics

Negative log-likelihood

Energy function

Score function

Force

Latent variables

Collective variables/coarse
graining/renormalization group

Partition function

Free energy calculation

Sample diversity

Enhanced sampling

Two sides of the same coin

Generative modeling



Known: samples

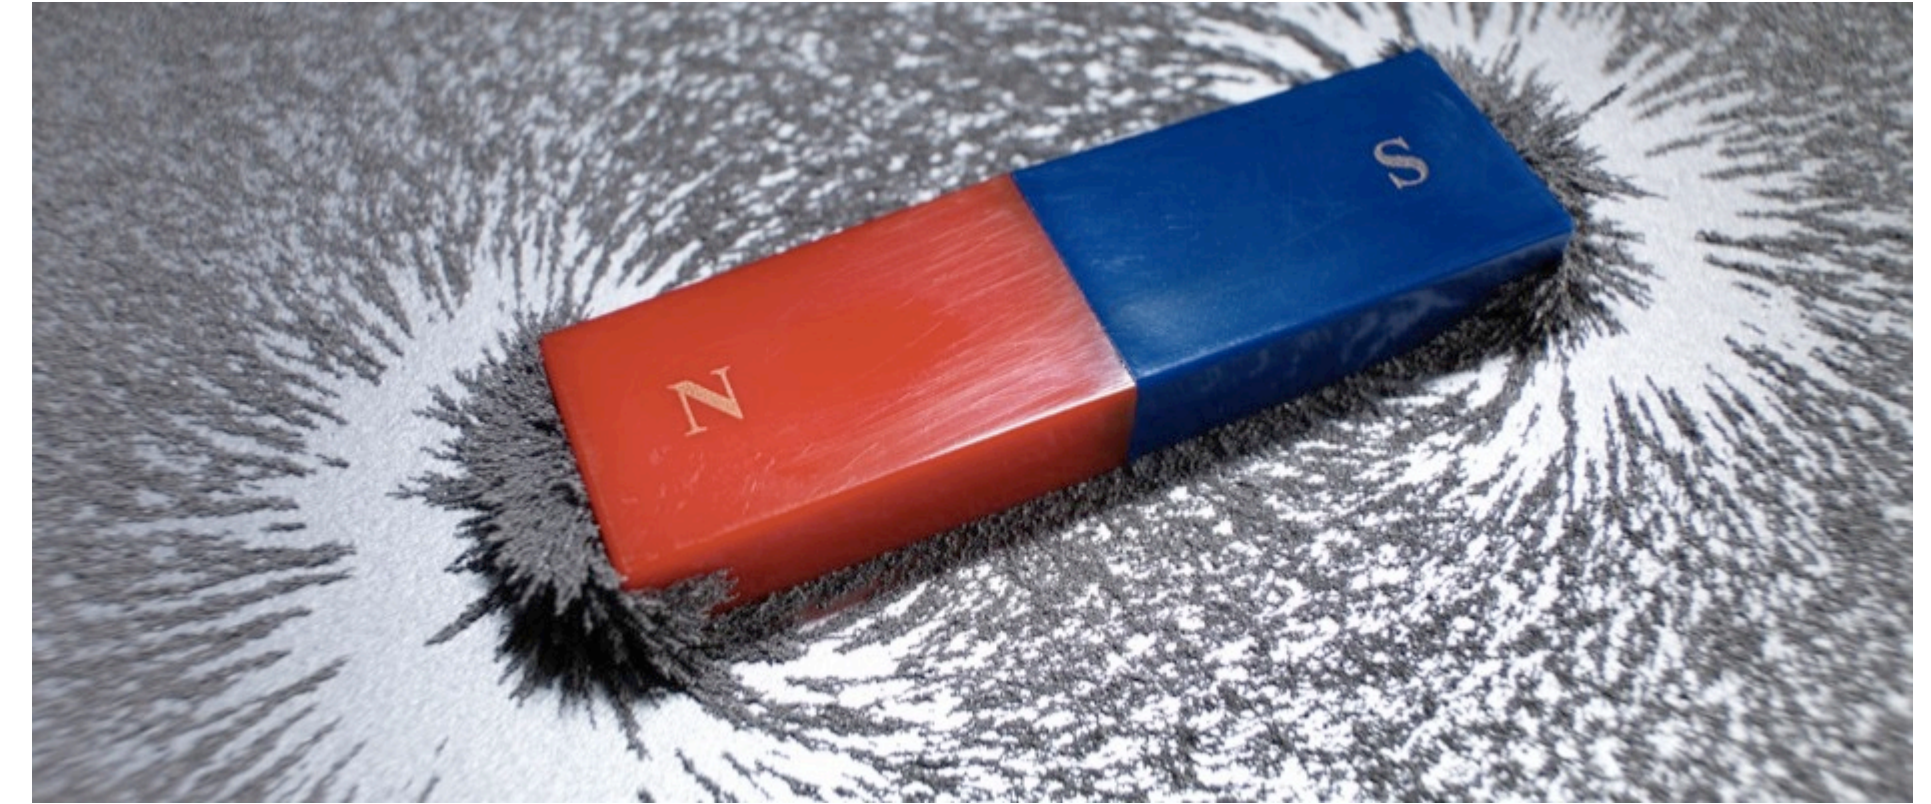
Unknown: generating distribution

Maximum likelihood estimation

“learn from data”

$$\mathcal{L} = - \mathbb{E}_{\mathbf{x} \sim \text{dataset}} [\ln p(\mathbf{x})]$$

Statistical physics



Known: energy function

Unknown: samples, partition function

Variational free energy

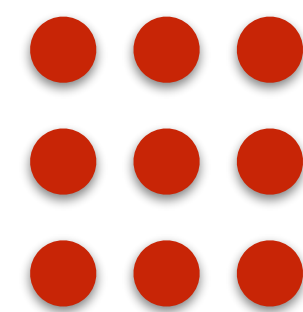
“learn from Hamiltonian”

$$F = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [H(\mathbf{x}) + k_B T \ln p(\mathbf{x})]$$

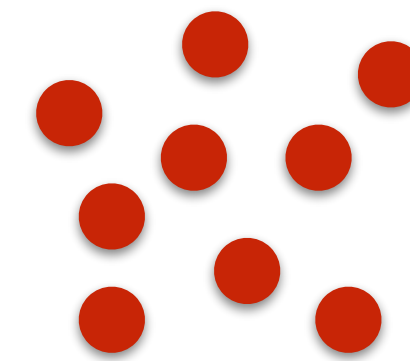
Nature tries to minimize free energy

$$F = E - TS$$

energy



entropy



F is the generating function of all other thermodynamic quantities

Unfortunately, it is “intractable” to compute

Deep variational free-energy approach

Use deep generative models as the variational density

$$F[p] = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \left[H(\mathbf{x}) + k_B T \ln p(\mathbf{x}) \right]$$

↓
energy

↓
entropy 😊

Li and LW, PRL '18
Wu, LW, Zhang, PRL '19

with normalizing flow &
autoregressive models



Tractable entropy



Direct sampling



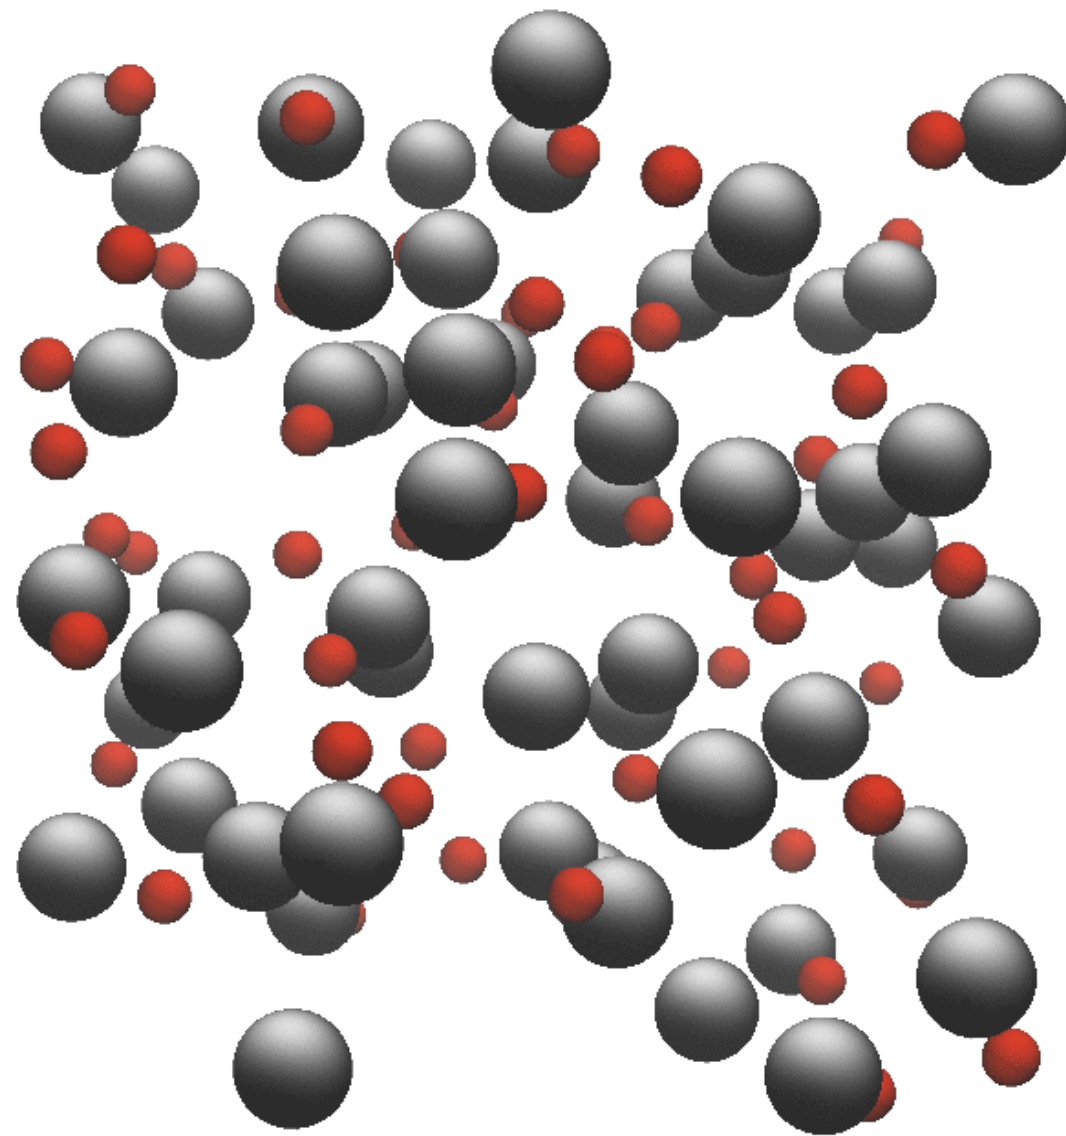
Turning a sampling problem to an optimization problem

better leverages the deep learning engine:



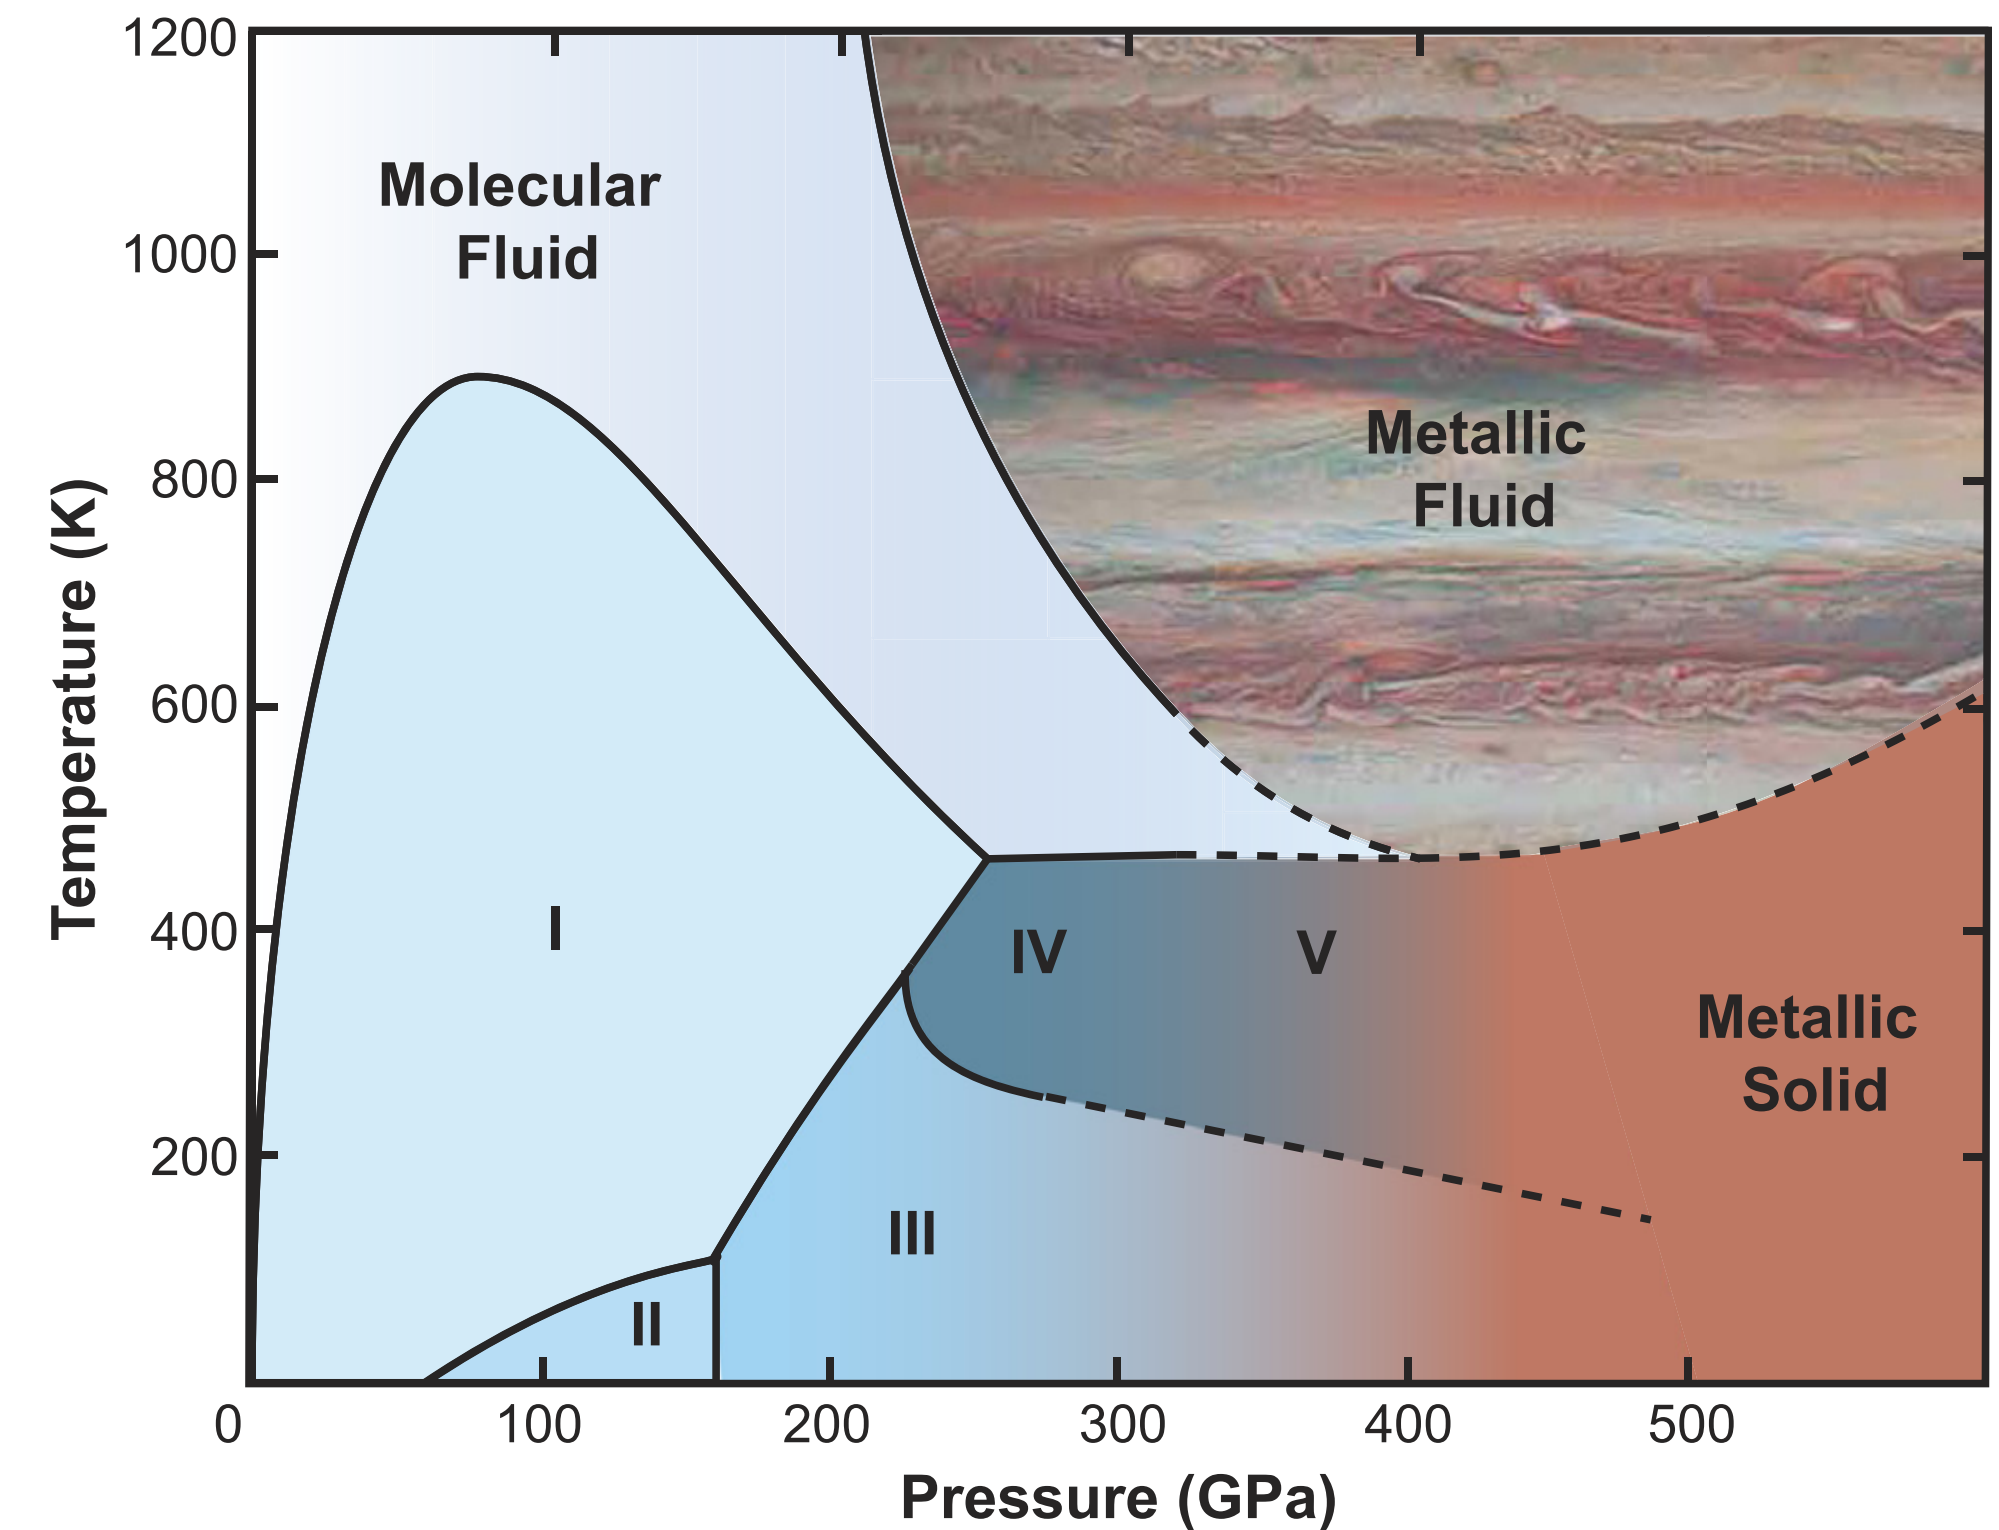
The dense hydrogen problem

Xie, Li, Wang, Zhang, LW, 2209.06095



N protons + N electrons in a box

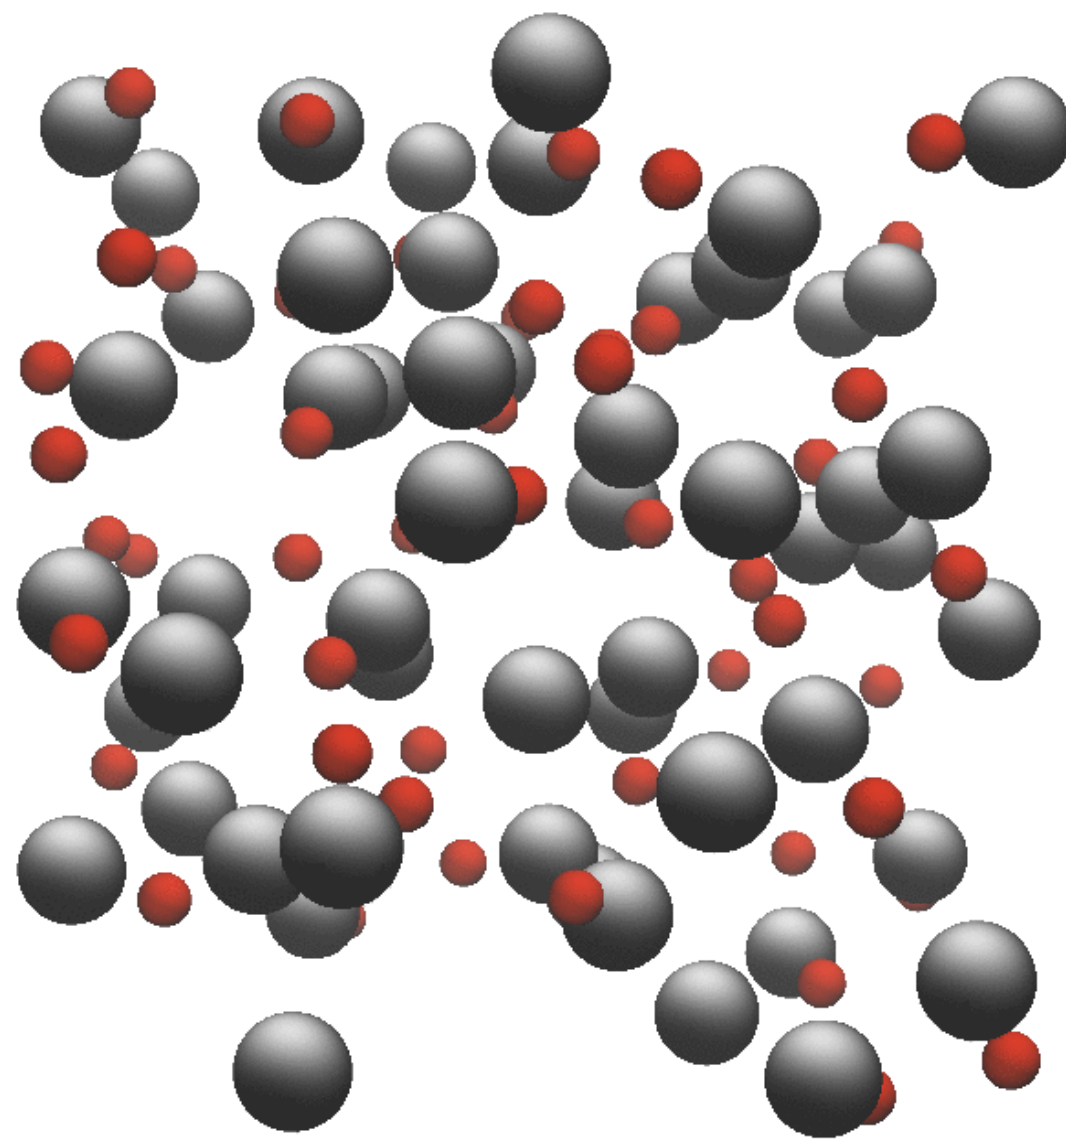
Radiat et al, Extremes, 2020



Generative model for proton probability density distribution
+ Deep neural network (Ferminet) for electron wavefunction

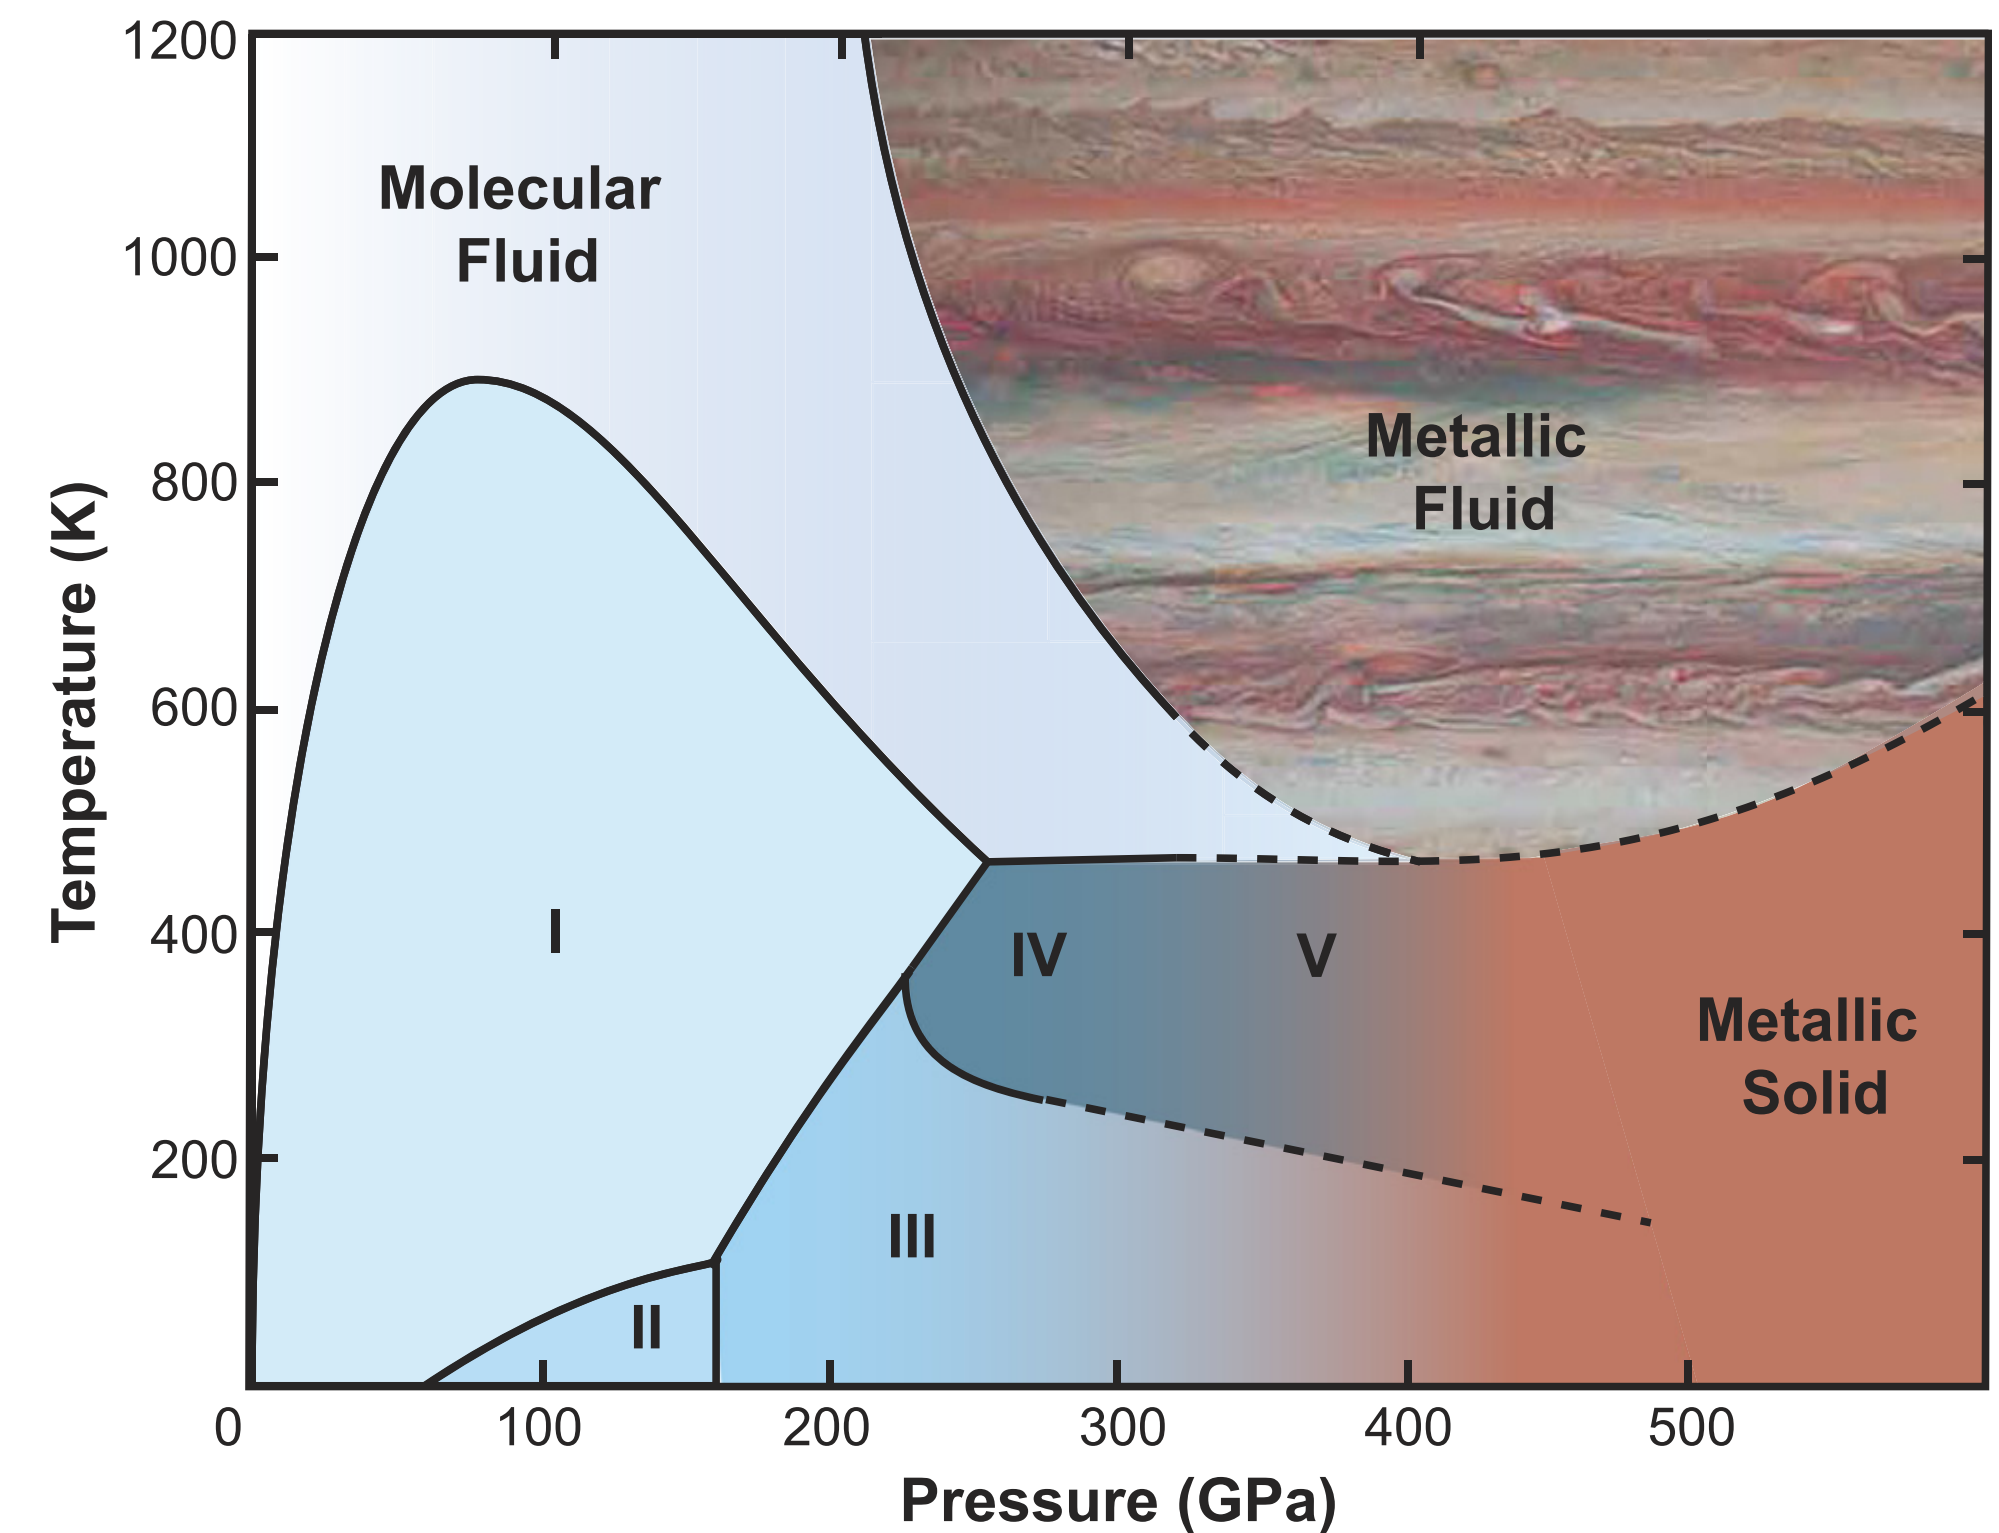
The dense hydrogen problem

Xie, Li, Wang, Zhang, LW, 2209.06095



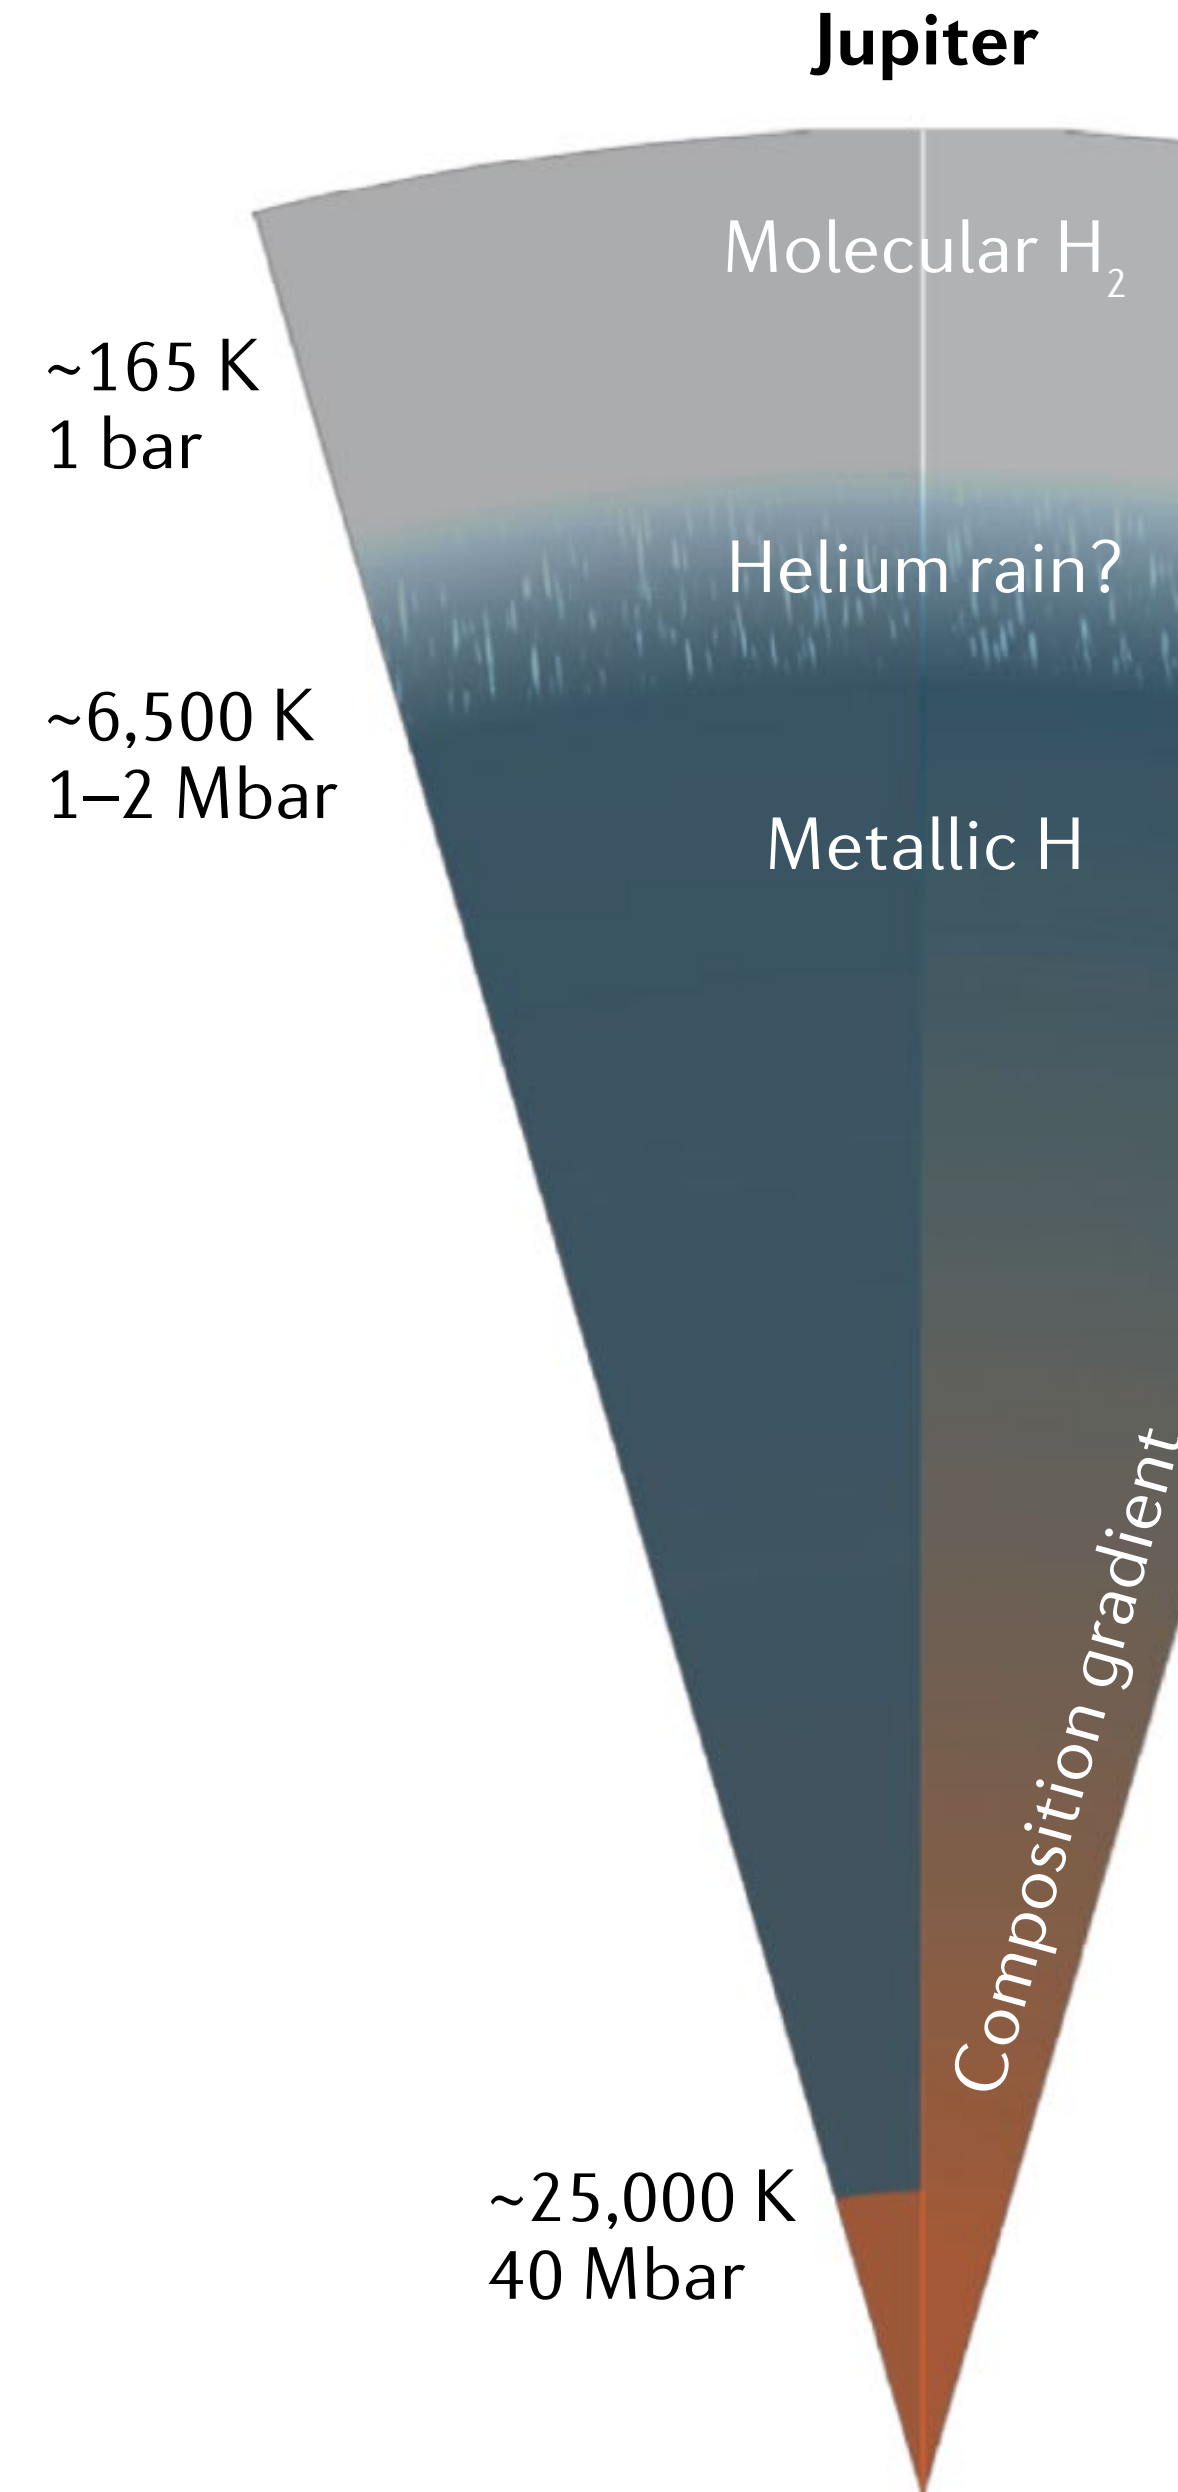
N protons + N electrons in a box

Radiat et al, Extremes, 2020

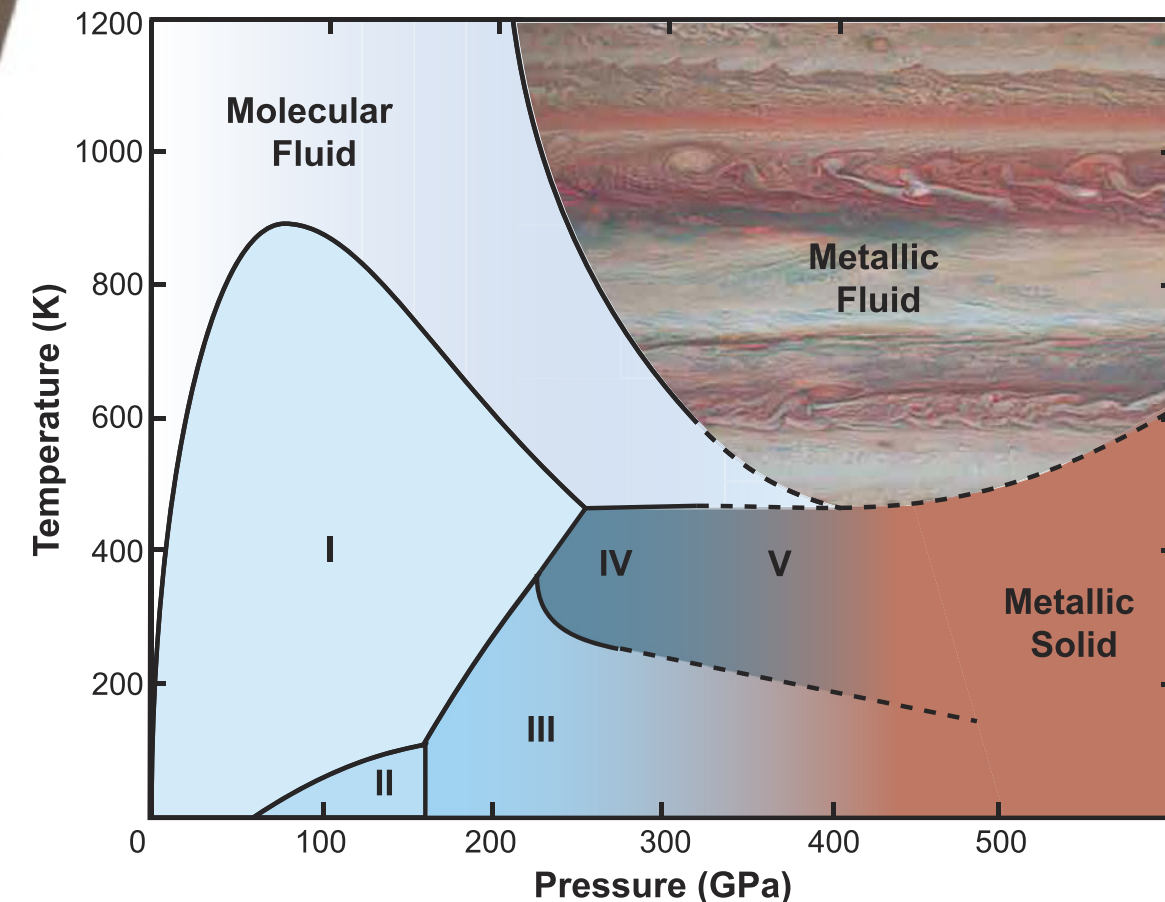


Generative model for proton probability density distribution
+ Deep neural network (Ferminet) for electron wavefunction

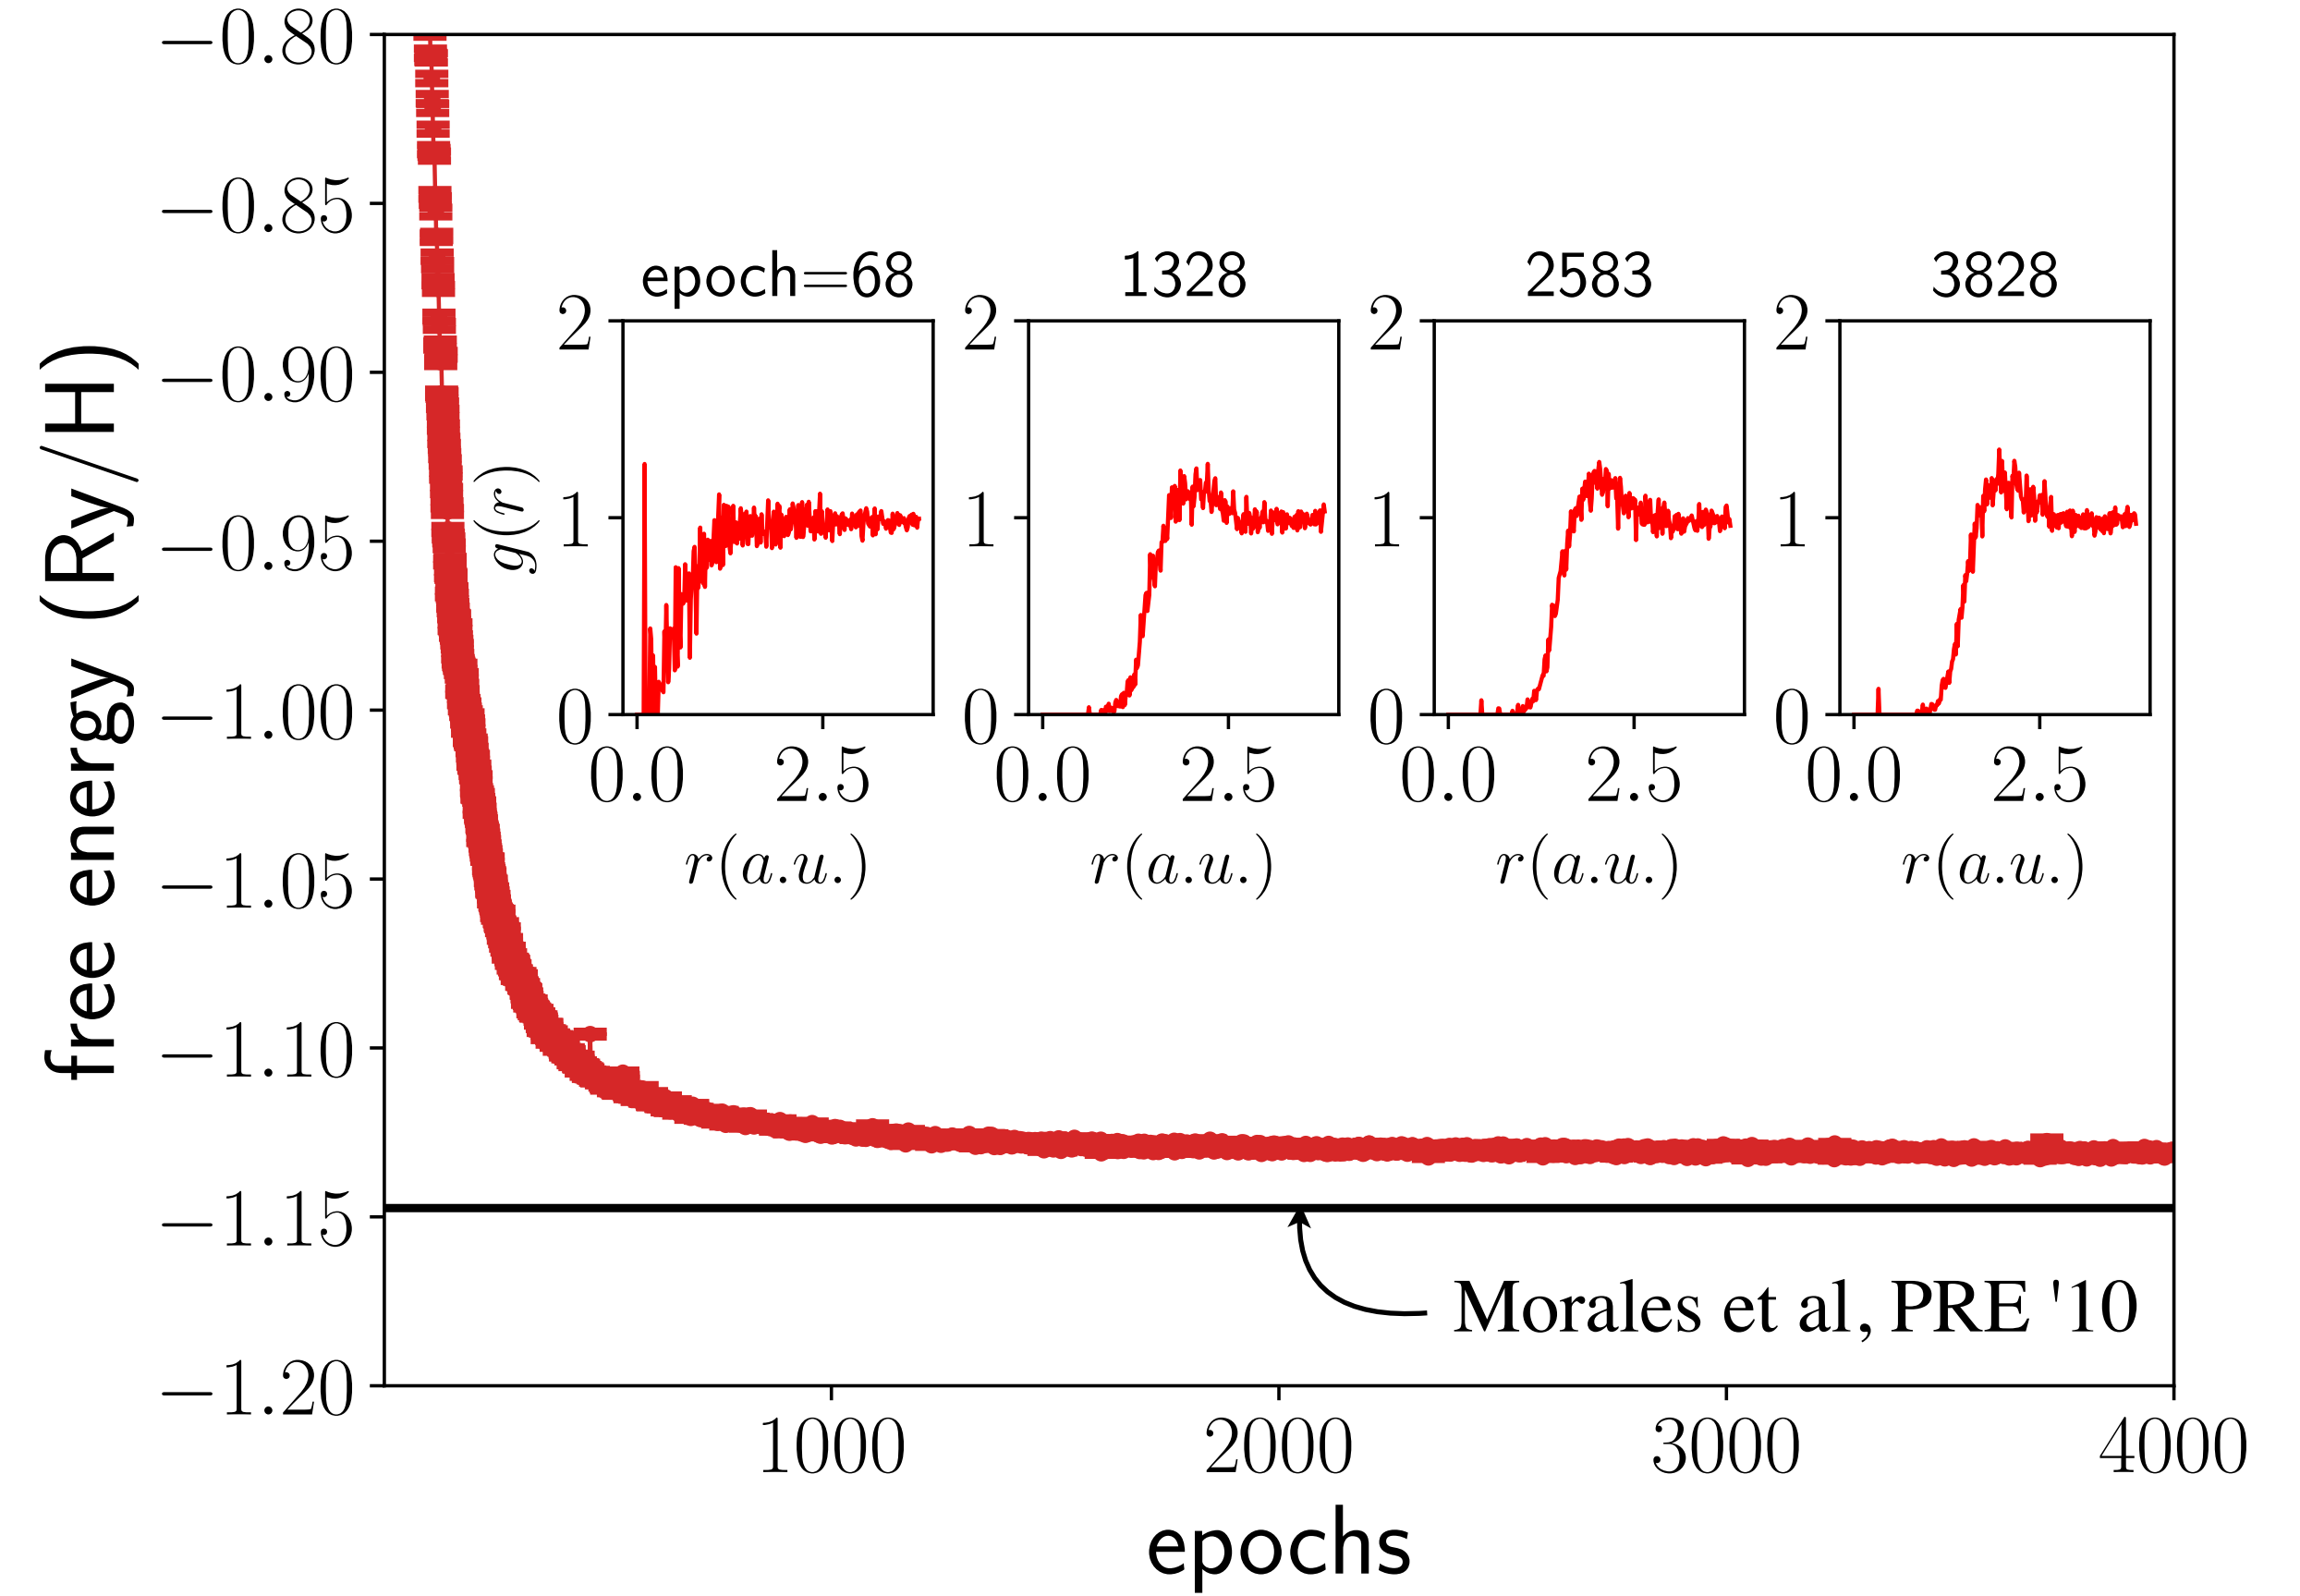
Dense hydrogen equation of state



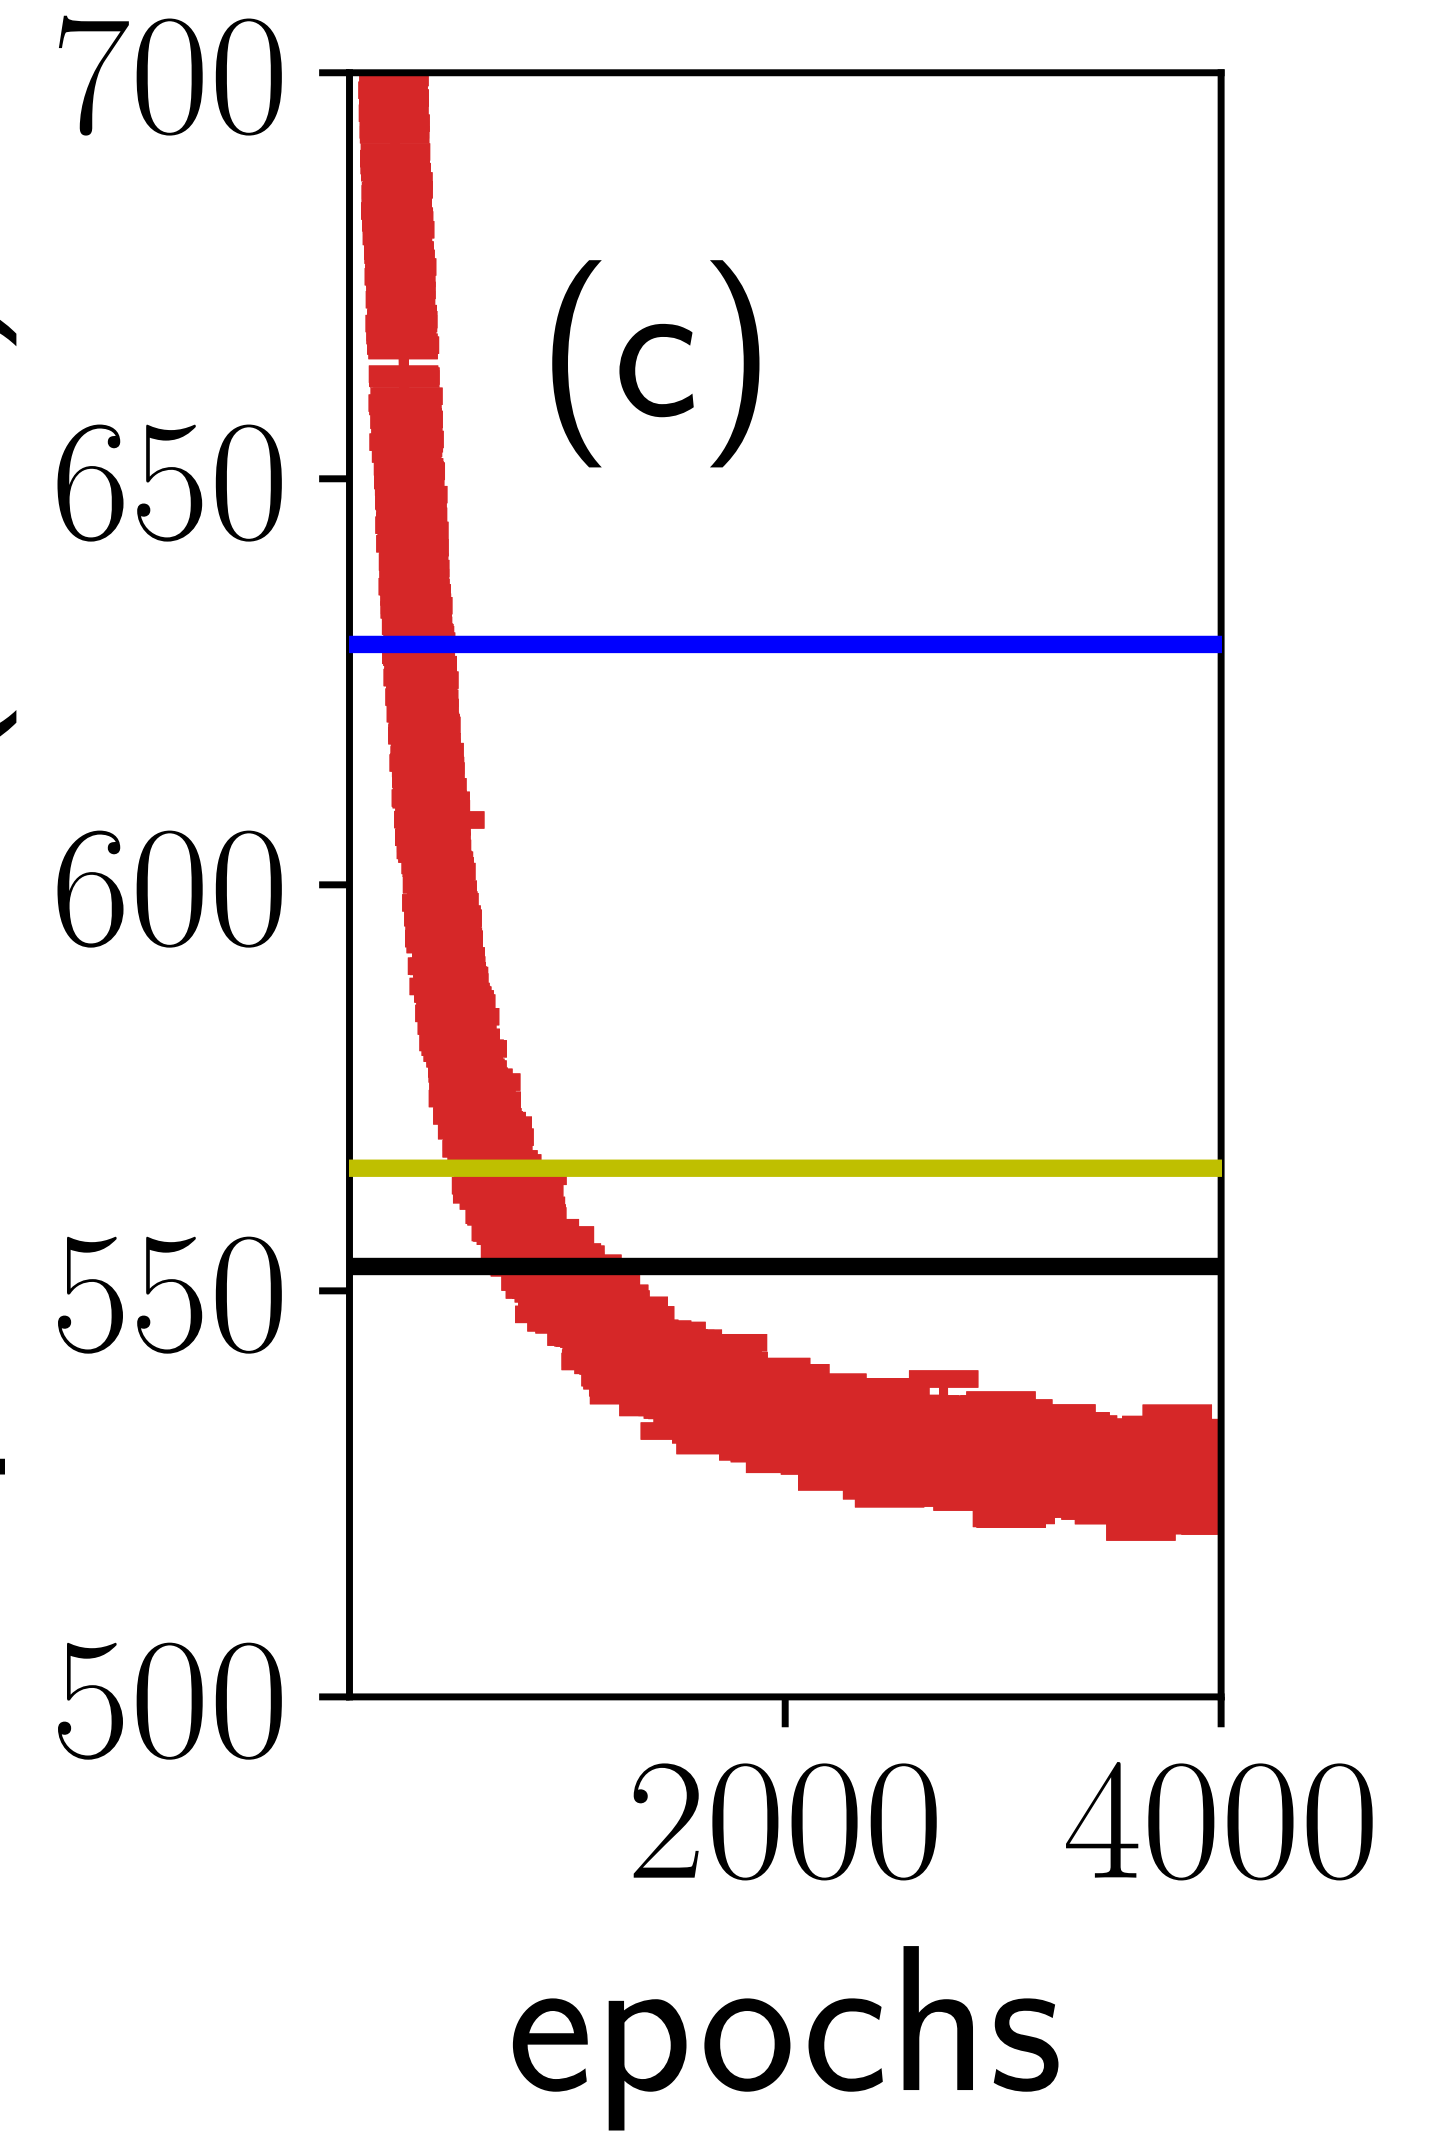
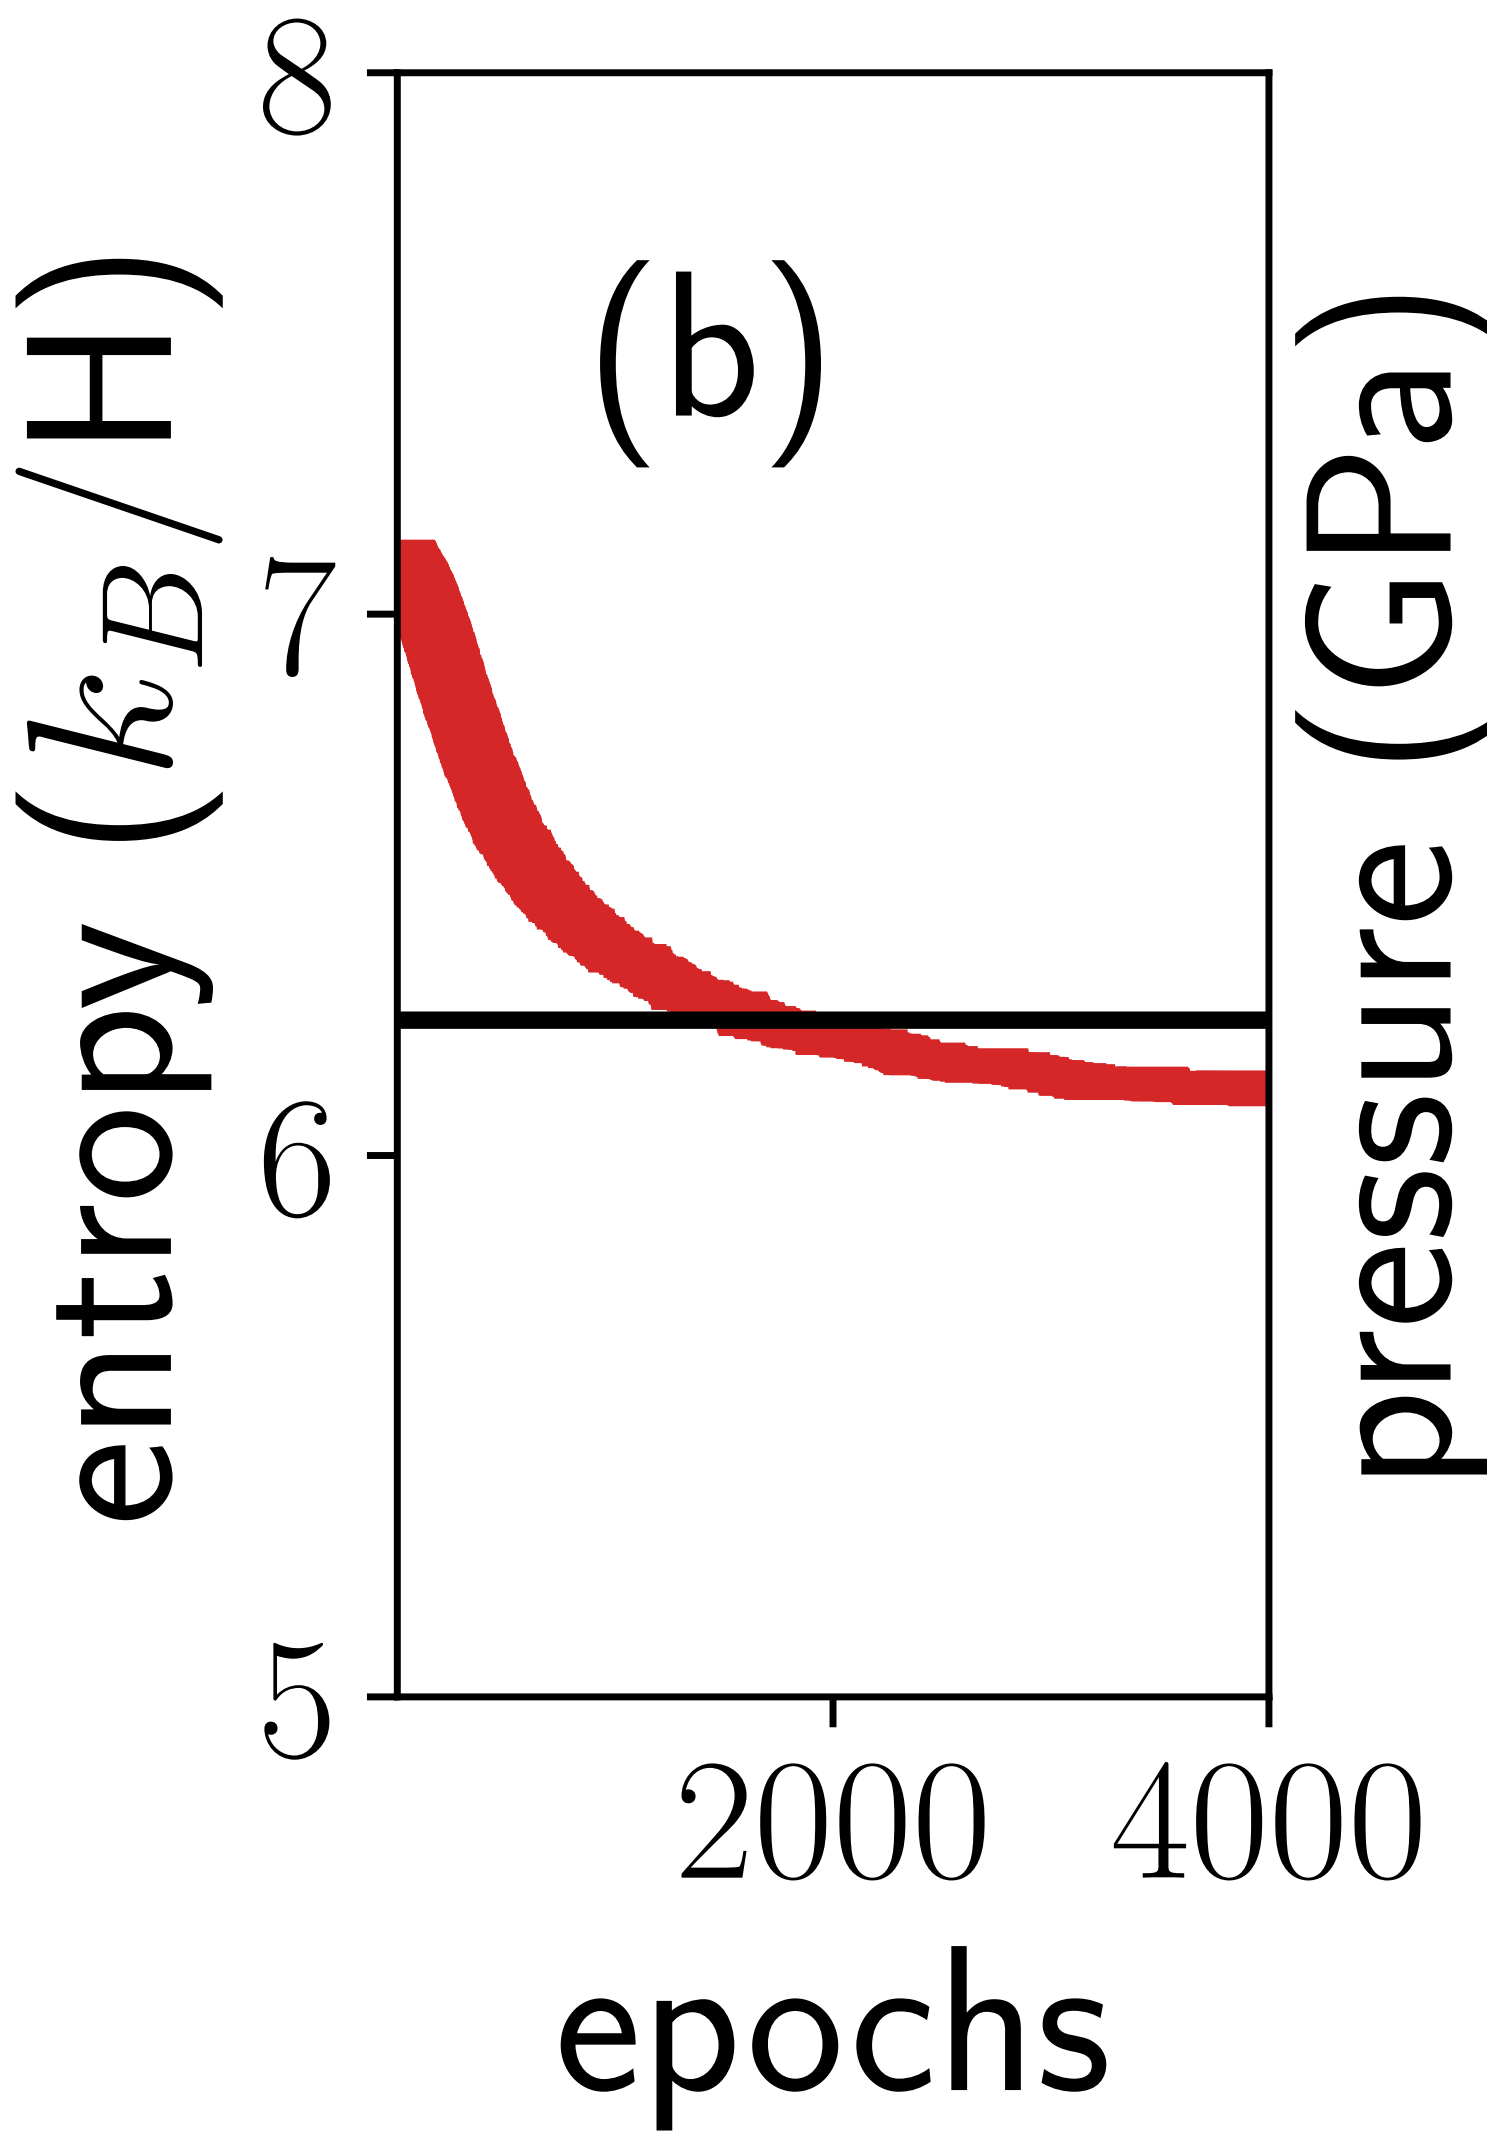
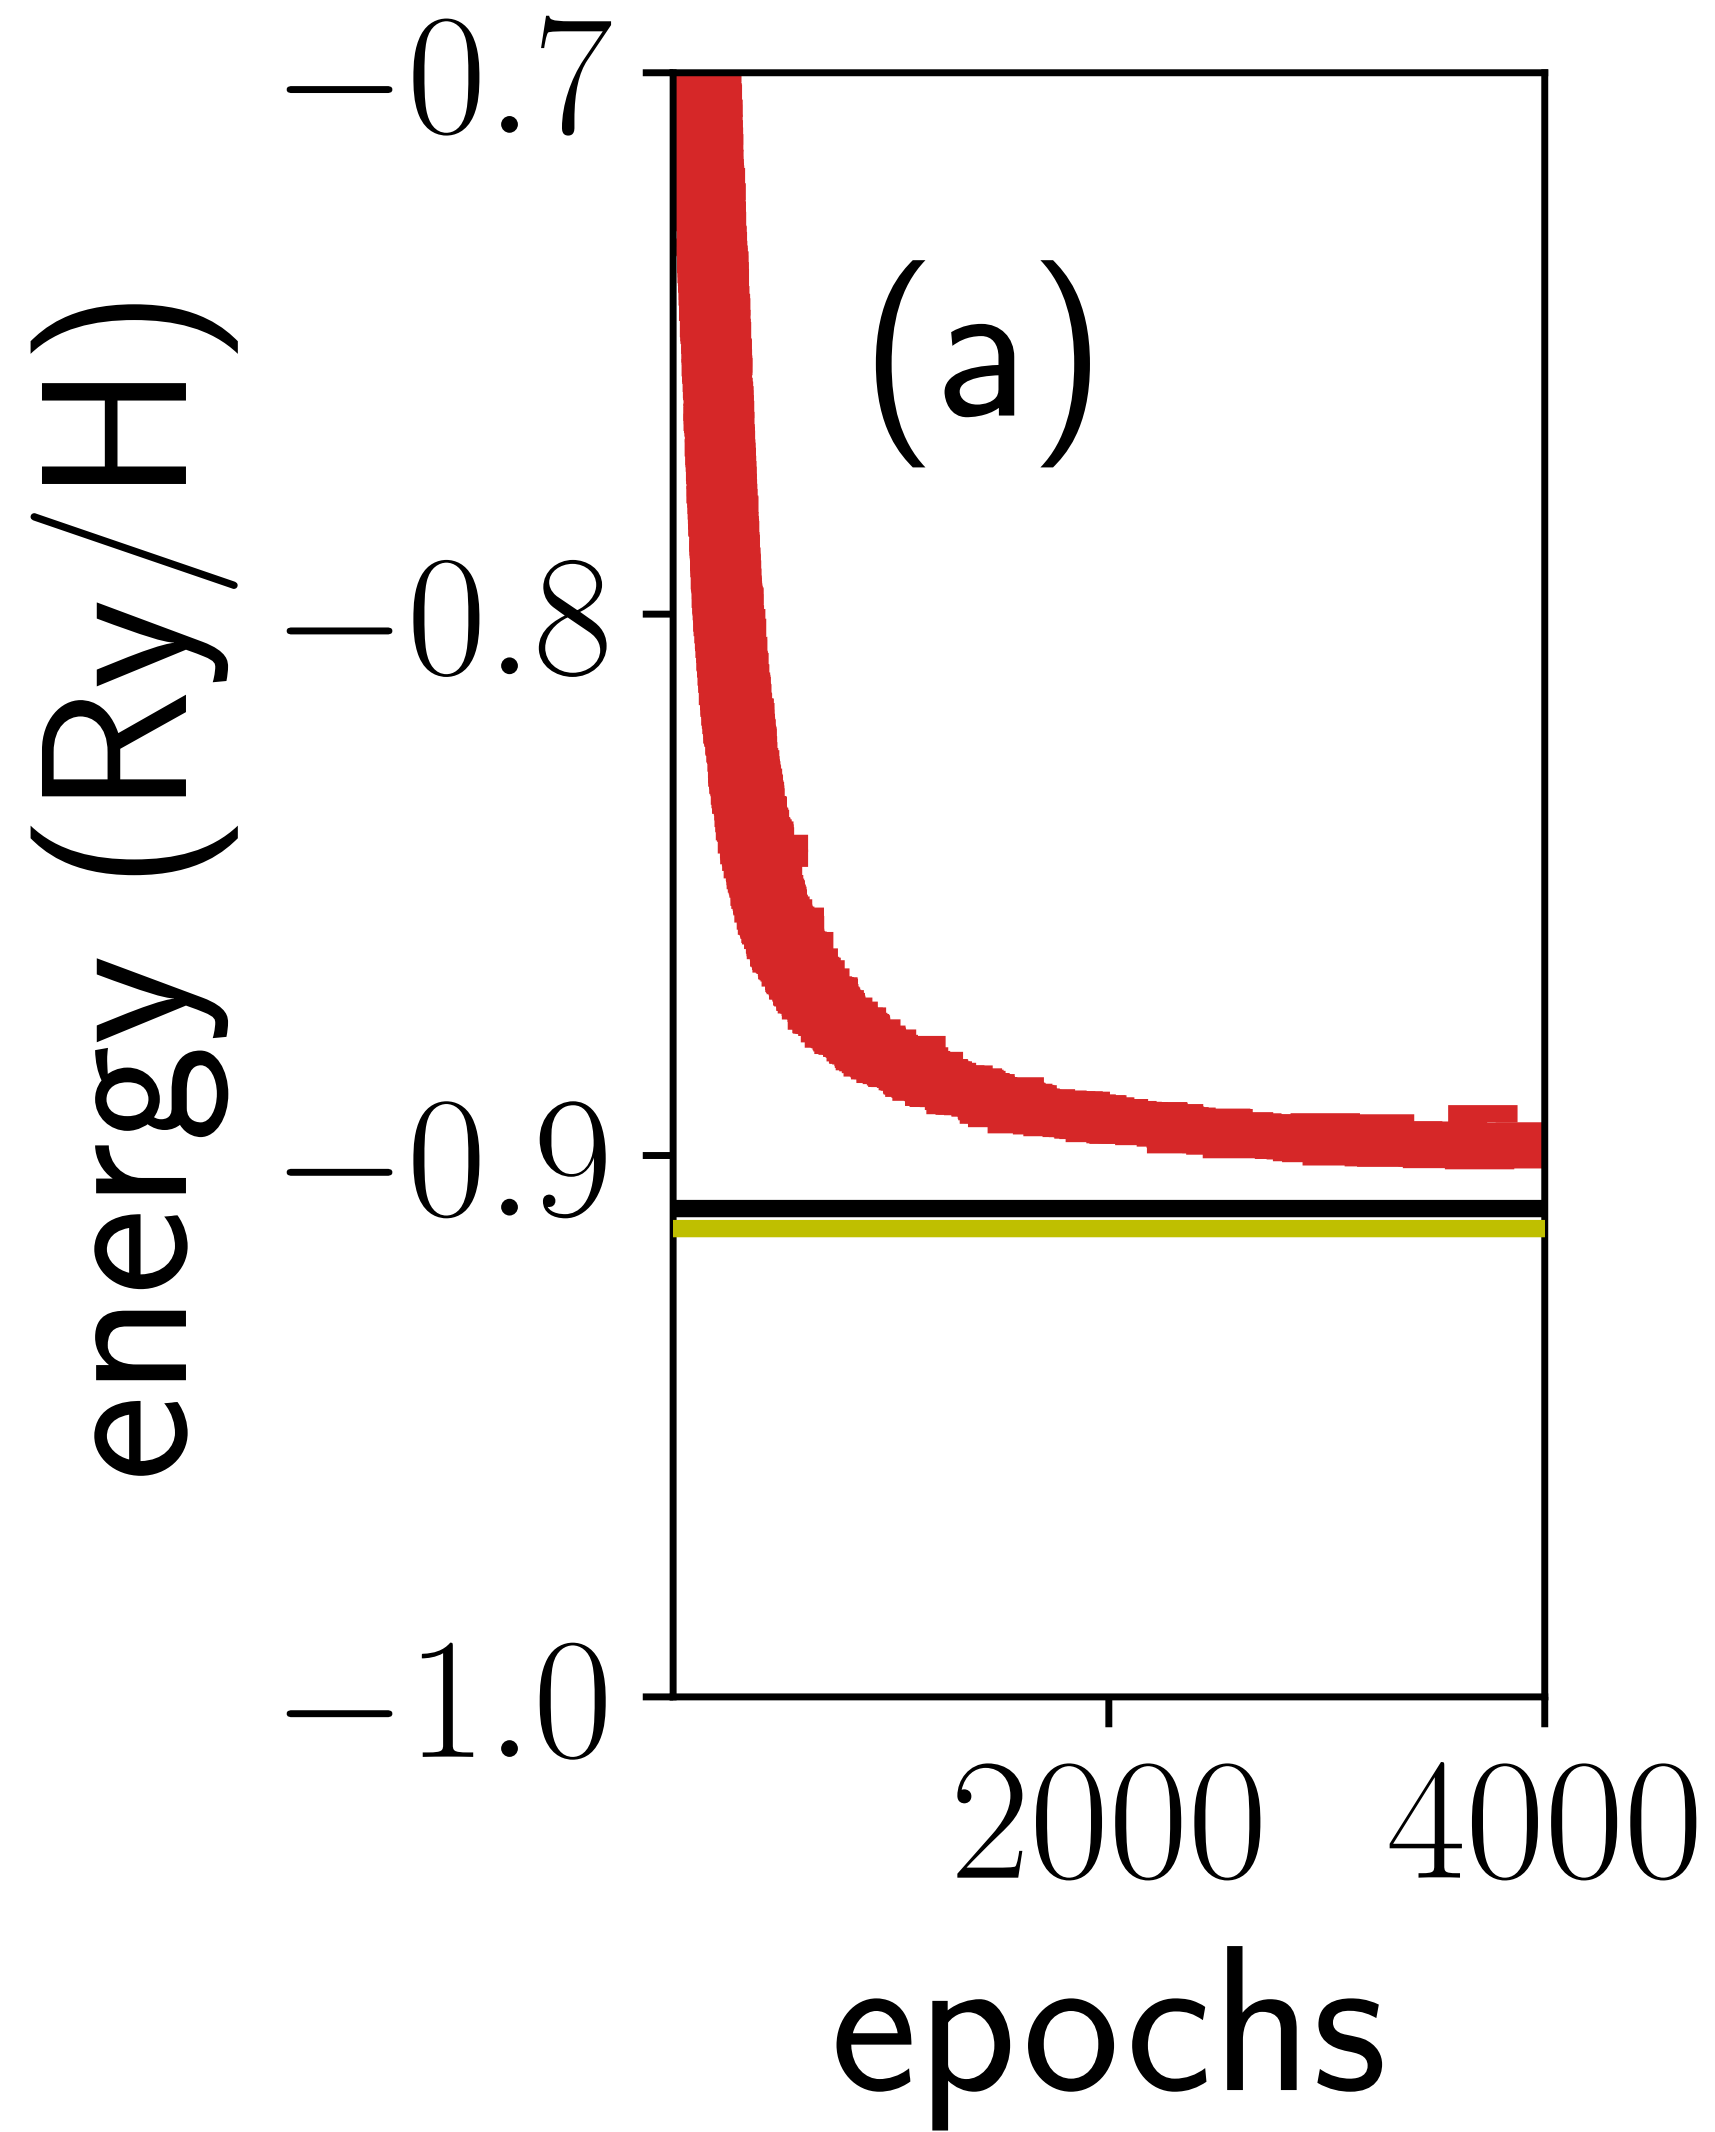
an adiabatic path in the phase diagram



$$r_s = 1.25 \quad T = 6000\text{K}$$



Xie, Li, Wang, Zhang, LW, 2209.06095



The Universe as a generative model

$$\mathcal{L} = \int d^4x \sqrt{-g} \left[\frac{m_p^2}{2} R - \frac{1}{4} F_{\mu\nu}^a F_a^{\mu\nu} + i \bar{\psi}^i \gamma^\mu D_\mu \psi^i + \left(\bar{\psi}_L^i V_{ij} \phi \psi_R^j + \text{h.c.} \right) - |D_\mu \phi|^2 - V(\phi) \right]$$



Thank you!

Discovering physical laws: **learning** the action
Solving physical problems: **optimizing** the action