

Interatomic Potentials Enabled by Machine Learning

Daniel Schwalbe-Koda

Lawrence Livermore National Laboratory
<https://dskoda.com>, <dskoda@llnl.gov>



Lecture: 4th IKZ-FAIRmat Winter School
Jan. 24, 2023

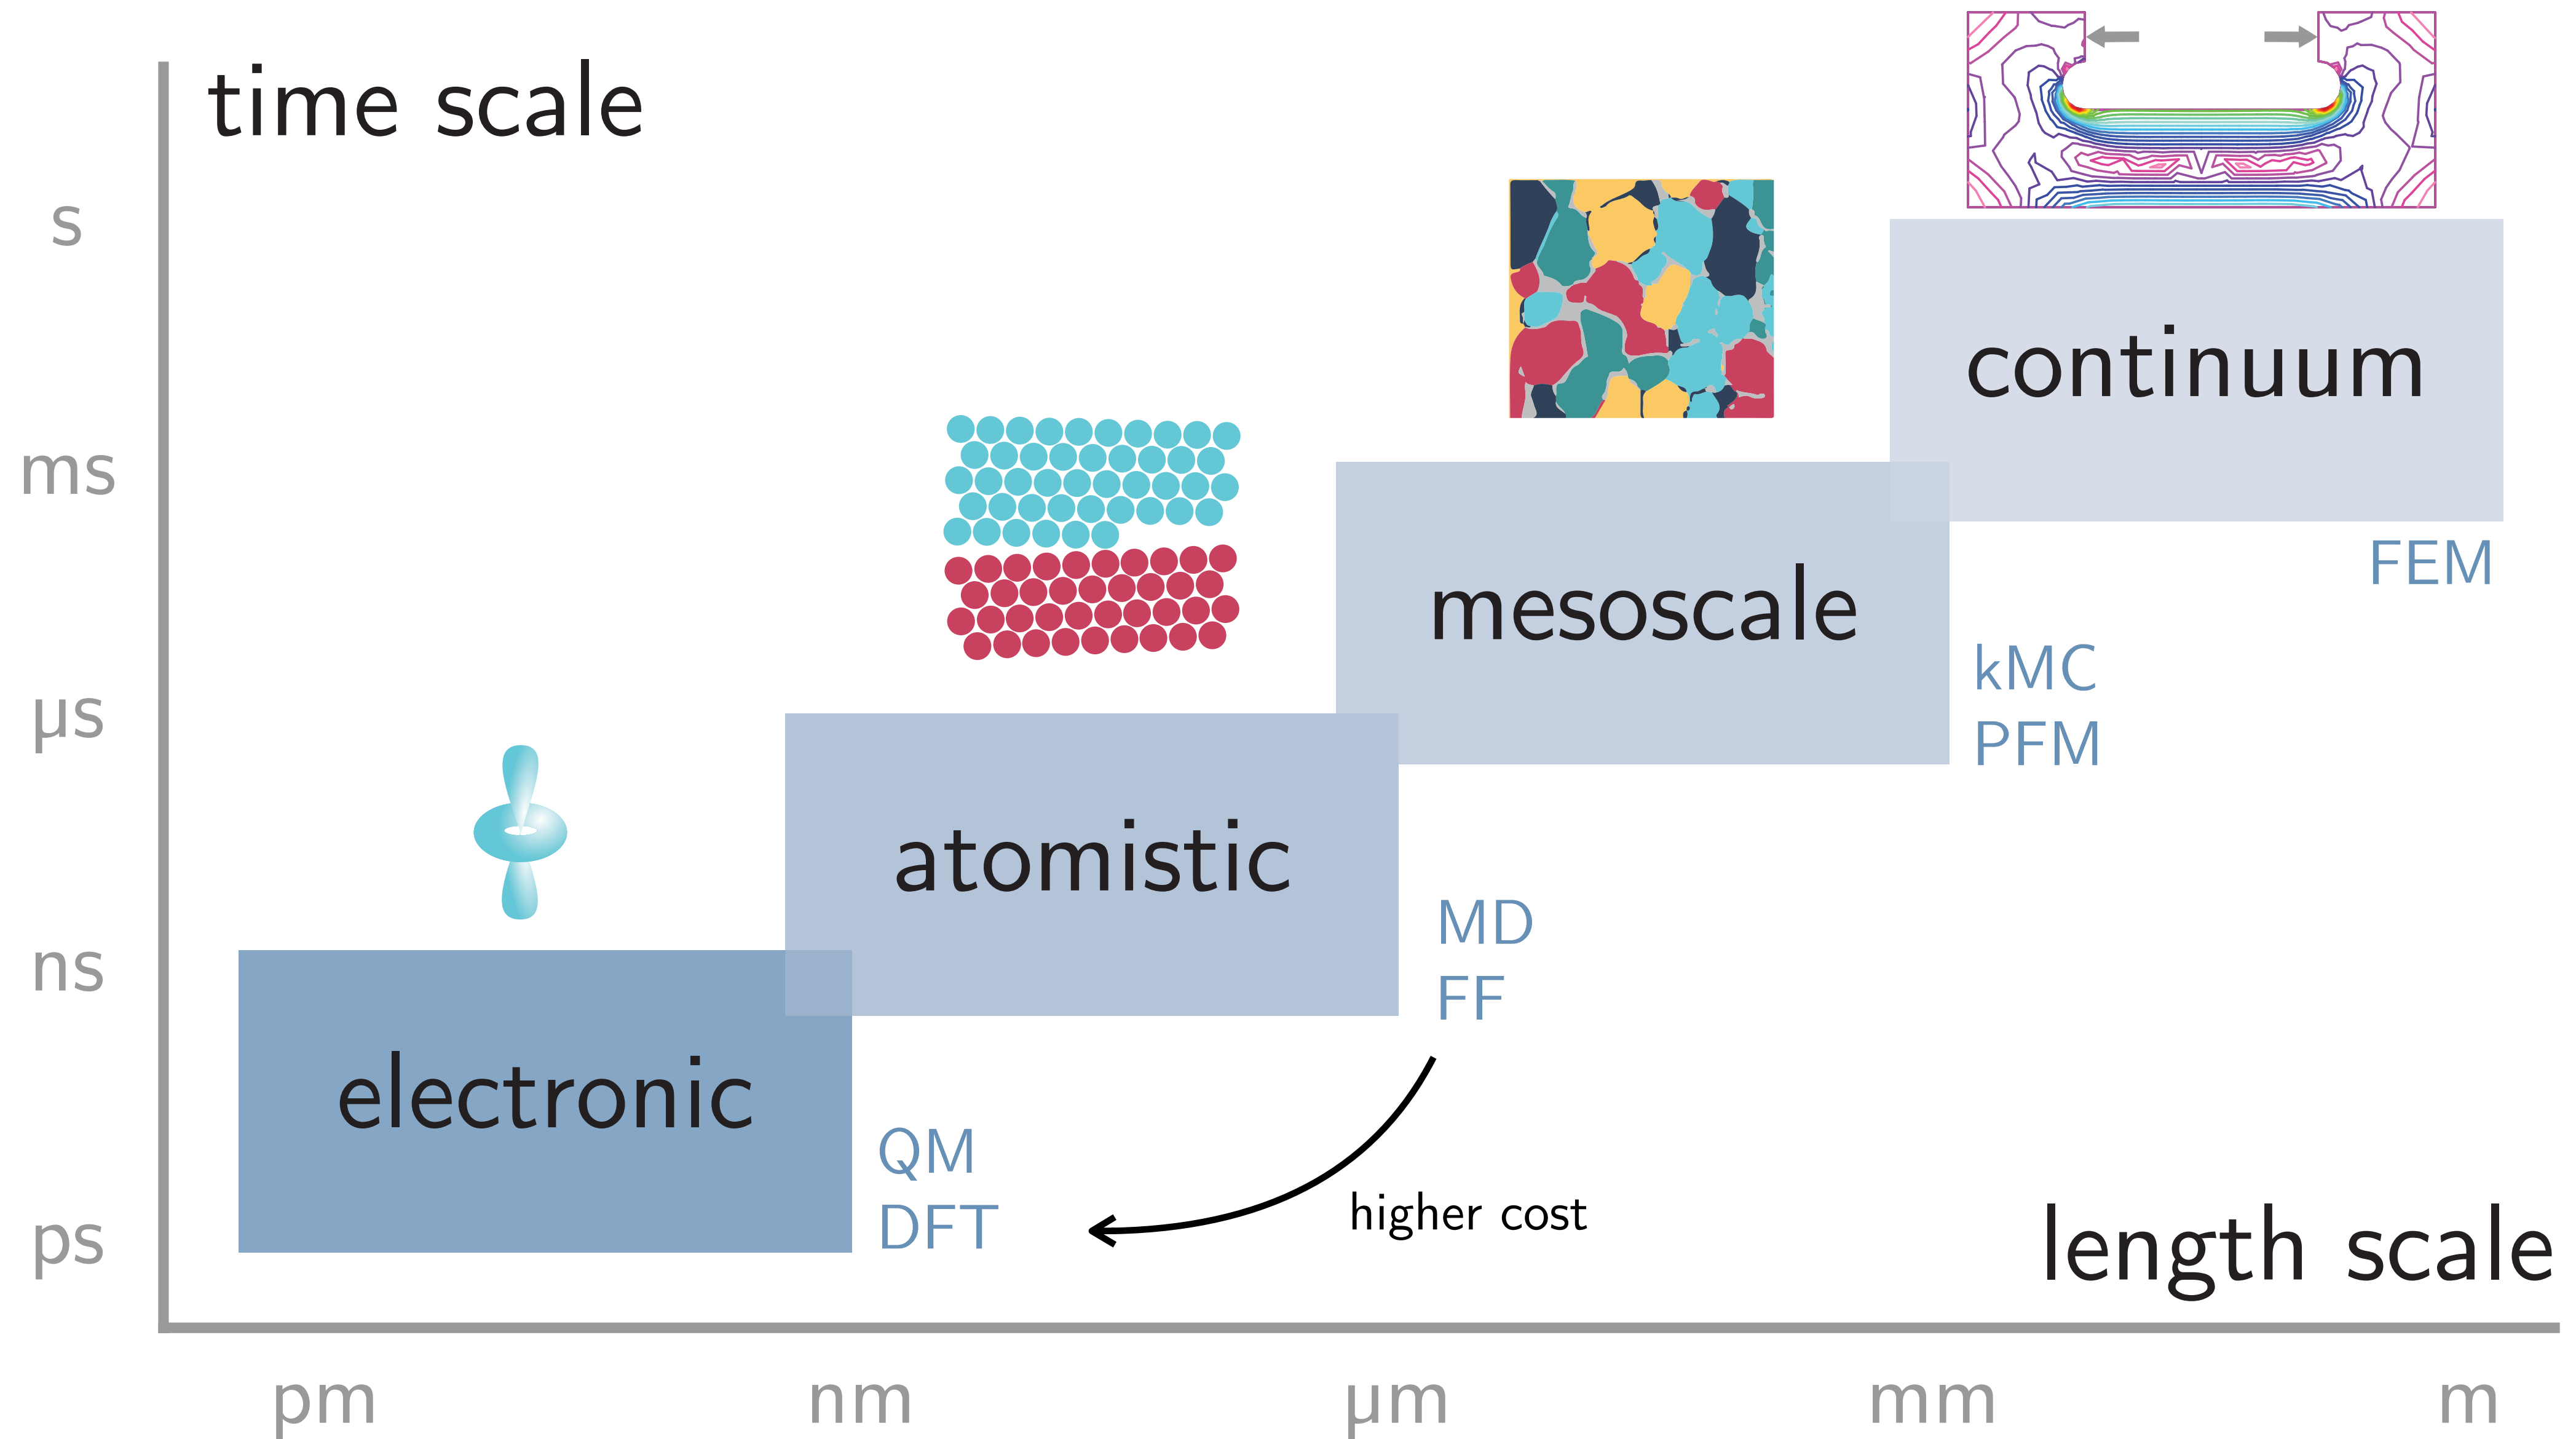
LLNL-PRES-843937

Prepared by LLNL under Contract DE-AC52-07NA27344. LDRD # 22-ERD-055

Image created with DALL-E 2

1. Introduction

Atomistic simulations and machine learning



In the atomistic world, several properties of interest are obtained from a **potential energy surface (PES)**

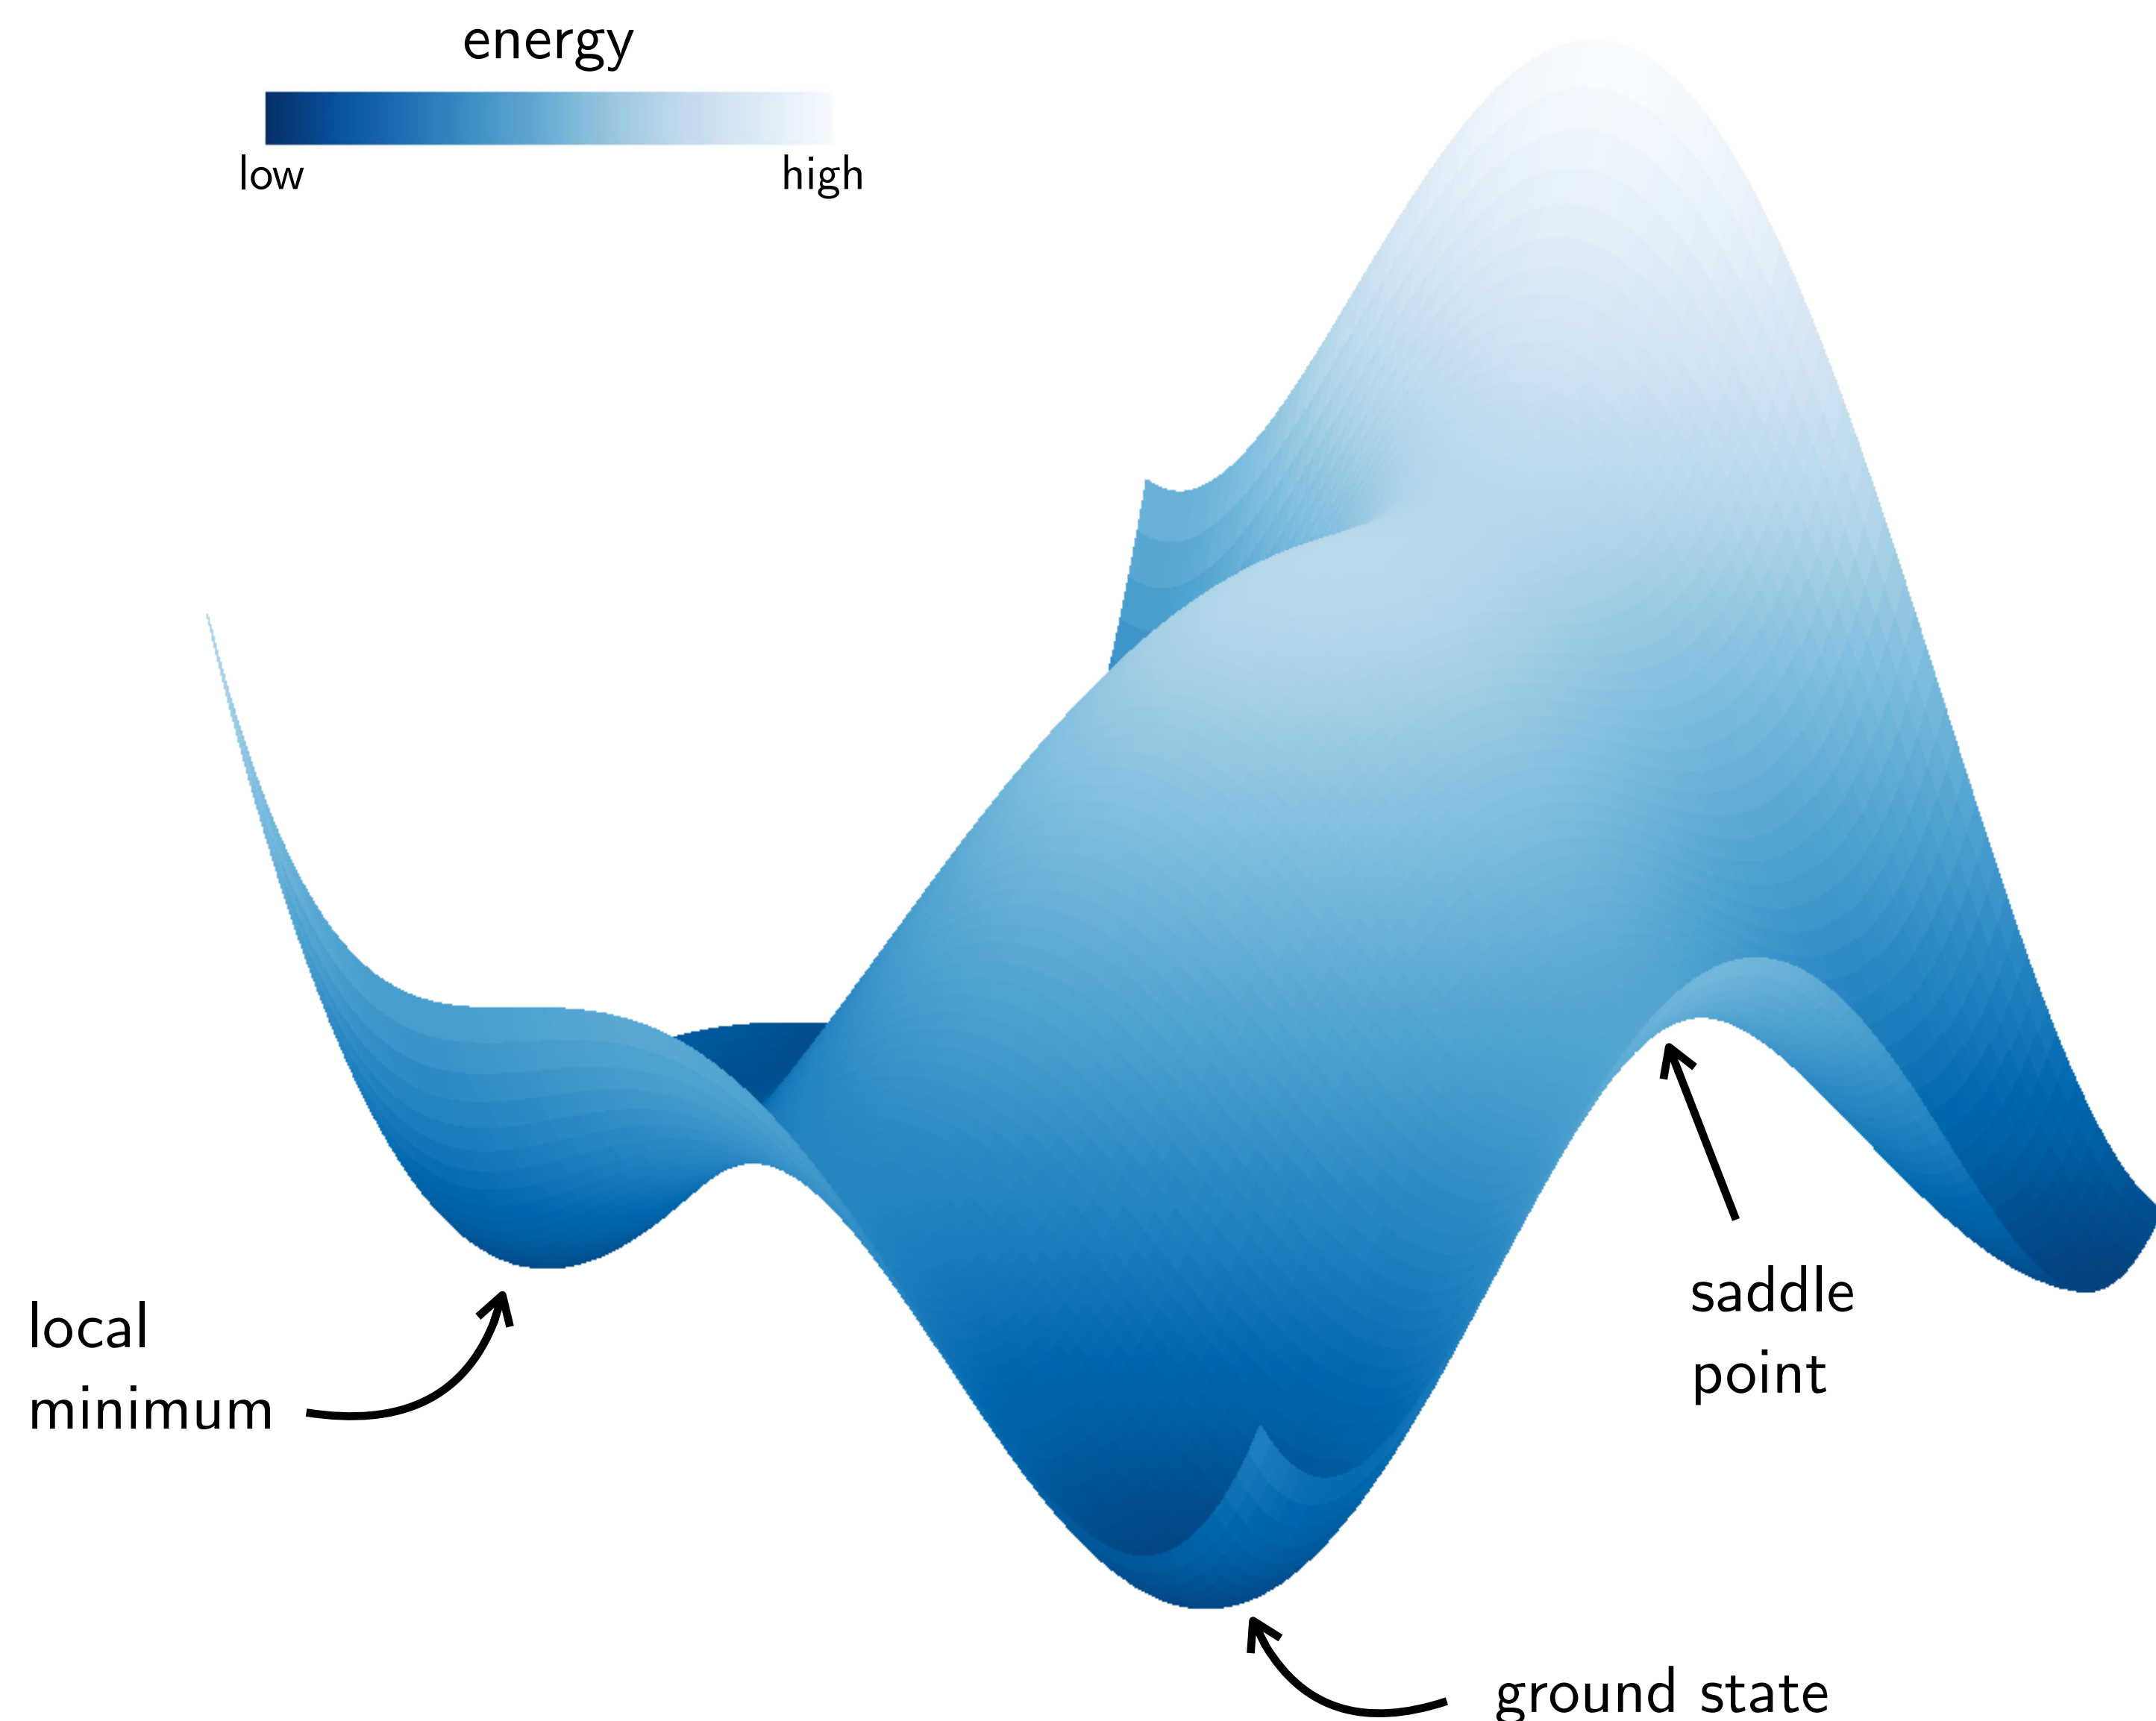
From a PES, we can derive: geometries, reaction energies, mechanical properties, dynamical behavior...

A PES is illustrated with respect to system coordinates (positions, distances, angles etc.)

However, in principle, the energy from atomistic systems can be computed from the atomic numbers and coordinates:

$$E = f(Z_i, \mathbf{r}_i)$$

(this is valid even in DFT or QC)



Sometimes, the properties of interest can be obtained from simpler models


Instead of using quantum mechanics or DFT to model a PES, we can use simpler models to describe interatomic interactions.

This reduces the computational cost: analytical energy functions are **much** faster than quantum mechanics calculations (several orders of magnitude).

For example, a pairwise potential is a simple approximation of the interaction energy between atoms:

$$E_{ij} = V(\mathbf{r}_i - \mathbf{r}_j)$$

Lennard-Jones potential: a simple example of pair potential



For example, potentials for noble gas usually rely on LJ models:

J. Chem. Phys., Vol. 119, No. 15, 15 October 2003





FIG. 1. Comparison of the Lennard-Jones potential used for rare gas interactions (labeled Ar-Ar) with the Mason-Schamp functions employed for Ar-K⁺ and Xe-Cs⁺. The appropriate Lennard-Jones pair well depth and pair equilibrium separation are taken as the units of energy and distance, respectively.

J. Hernandez-Rojas, D. Wales. *JCP* **119** (15), 7800 (2003)

Lennard-Jones potential: a simple example of pair potential



However, for more complicated systems, there are several options of potentials such as:

Lennard-Jones potential

$$V(r) = 4\epsilon \left[\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right]$$

Morse potential


$$V(r) = A \left[1 - e^{-a(r-r_{\text{eq}})} \right]^2$$

Harmonic potential

$$V(r) = A(r - r_{\text{eq}})^2$$

Buckingham-Coulomb potential

$$V(r) = Ae^{-Br} - \frac{C}{r^6} + \frac{q_1q_2}{4\pi\epsilon_0r}$$



Usually, interatomic potentials involve a combination of these terms

e.g., the CHARMM22 force field:

$$\begin{aligned}
 V(r) = & \sum_{\text{bonds}} k_b (b - b_0)^2 + \sum_{\text{angles}} k_\theta (\theta - \theta_0)^2 + \sum_{\text{improp.}} k_\omega (\omega - \omega_0)^2 + \sum_{\text{Urey-Bradley}} k_u (u - u_0)^2 \\
 & + \sum_{\text{dihedrals}} k_\phi [1 + \cos(n\phi - \delta)] + \sum_{\text{nonbond.}} 4\epsilon_{ij} \left[\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^6 \right] + \sum_{\text{nonbond.}} \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}}
 \end{aligned}$$

Classical force fields are usually fit to structural, vibrational, and other energy-based models from *ab initio* calculations.

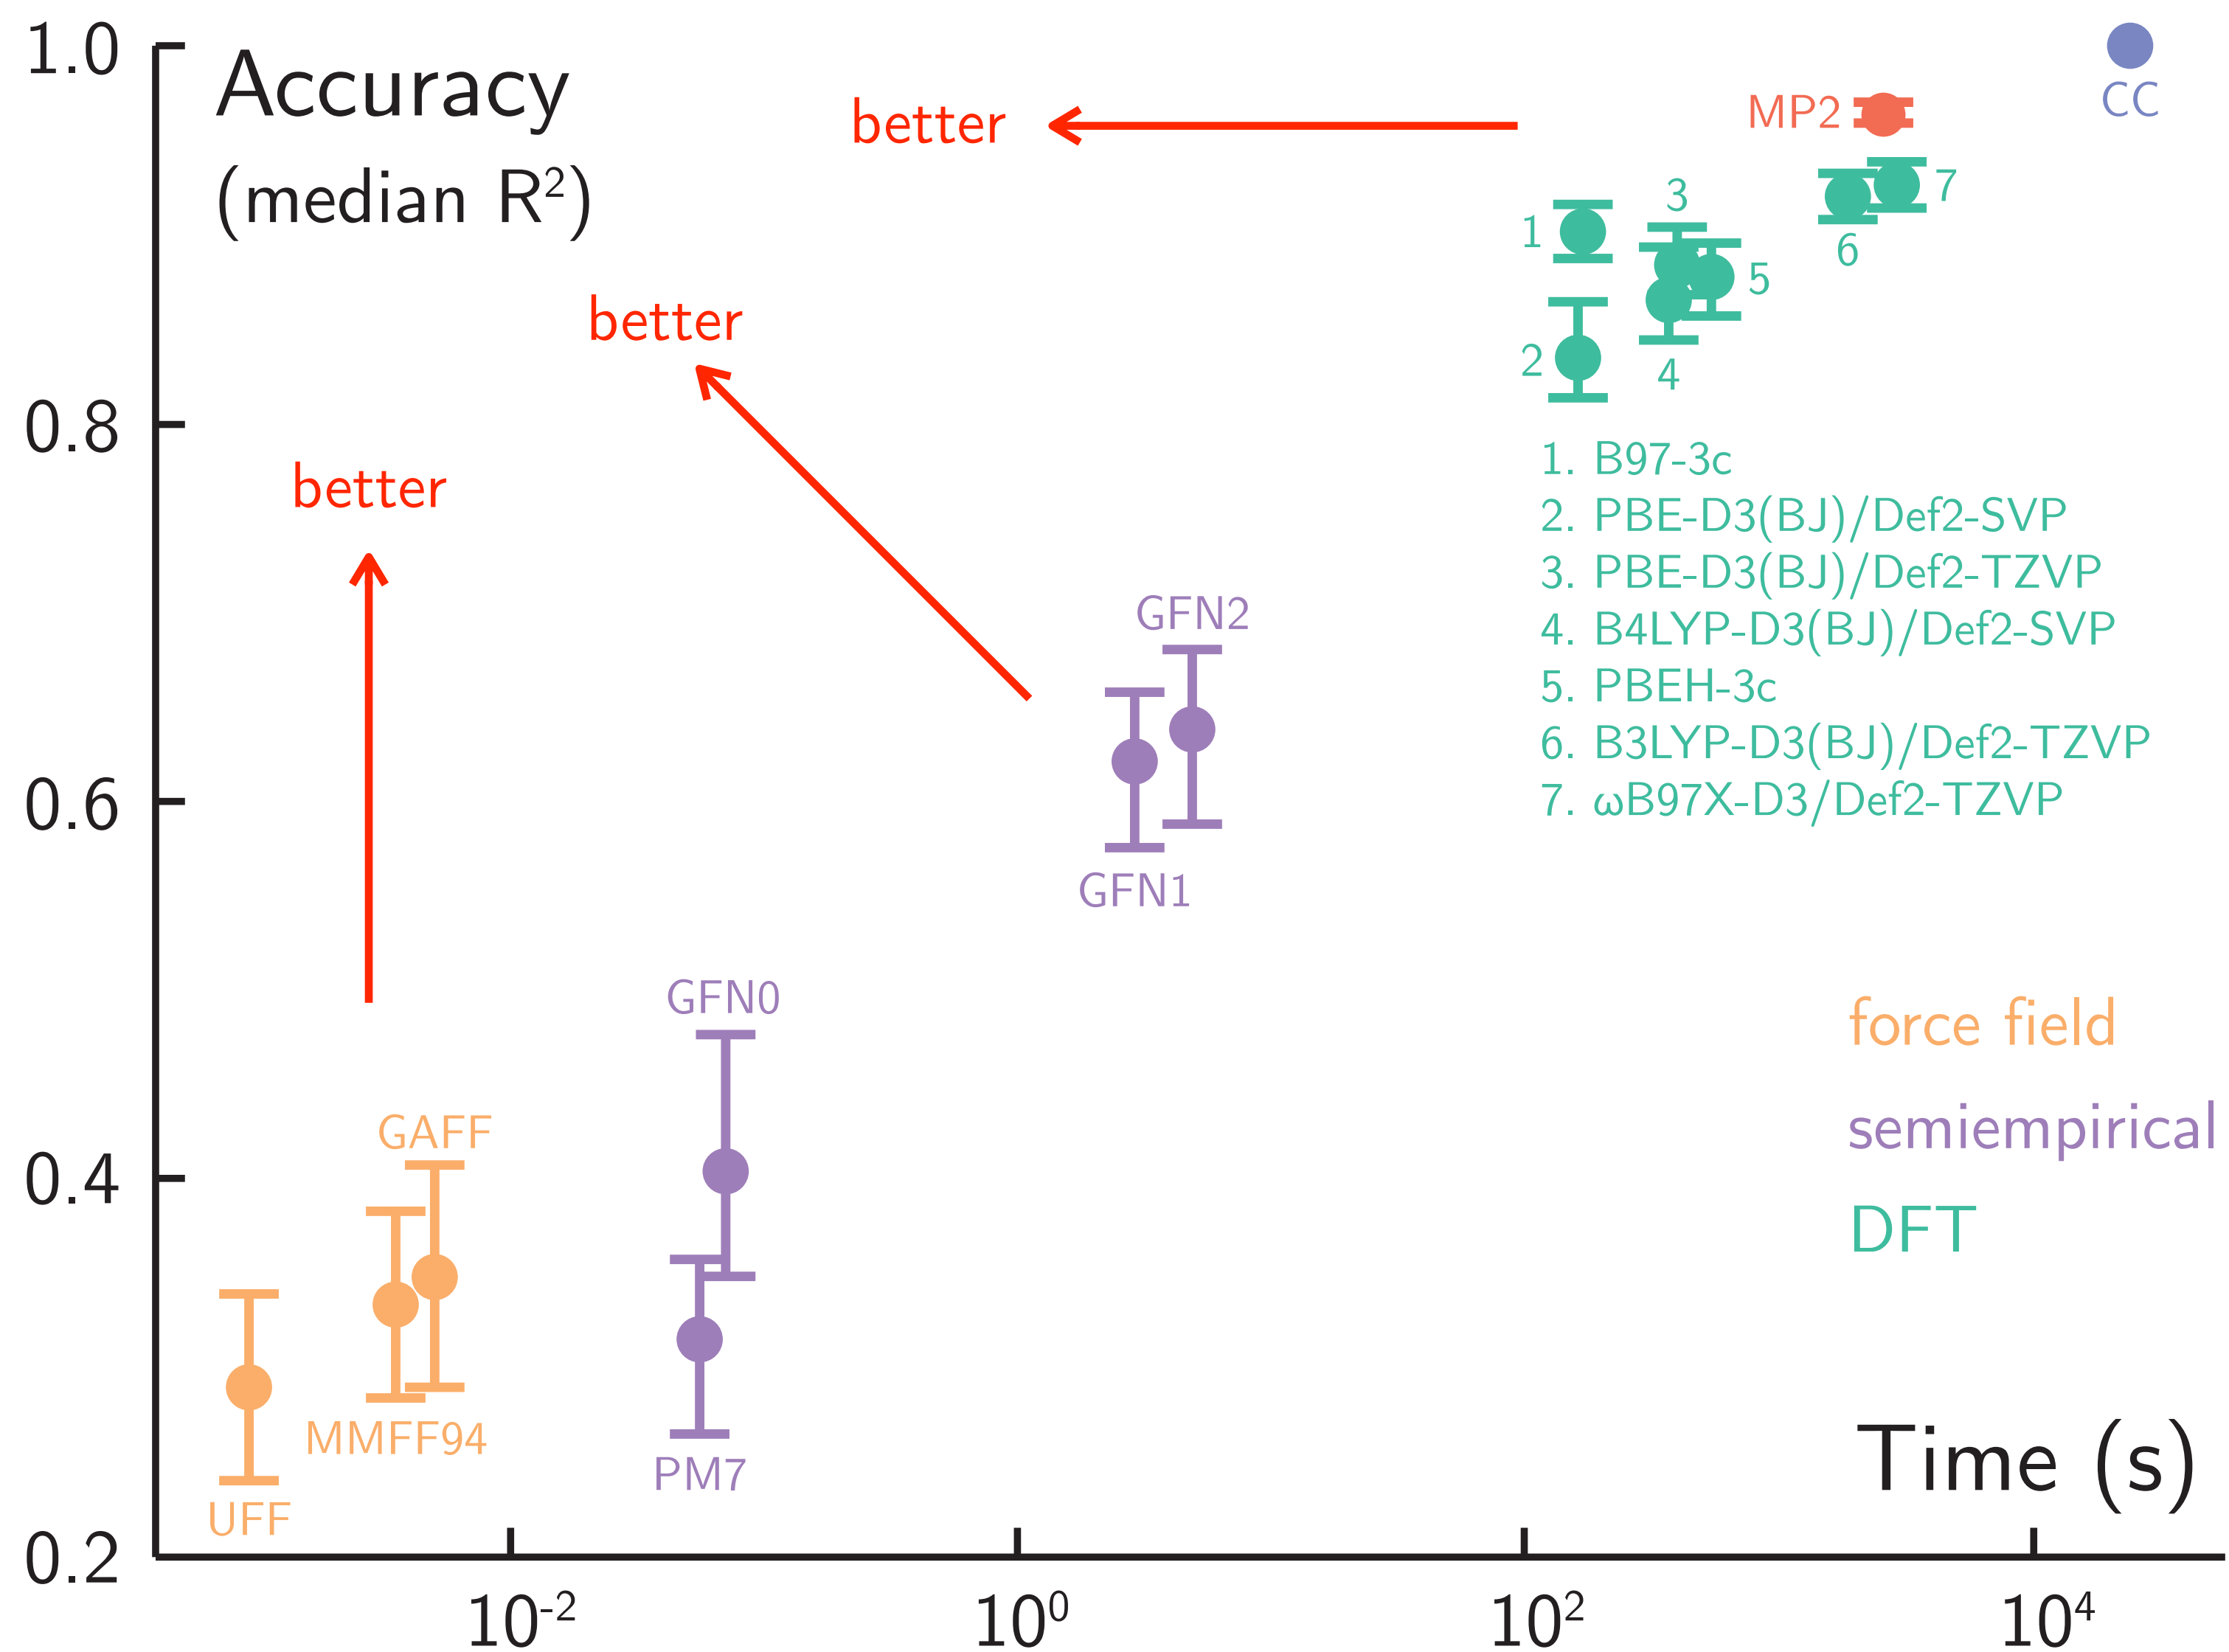
The problem is: how to choose the functional forms and parameters?

From the CHARMM22 paper:

Adjustment of the parameters was performed manually, although in certain cases (e.g., for proline) automated procedures were employed. We have found that automated procedures must be used with great care owing to the extensive nature of parameter space, correlation among the parameters, and their underdetermined nature. An automated least-squares procedure often leads to a combination of “unphysical” parameters that reproduce the input data. More meaningful parameter values, which have a wider range of applicability, were obtained manually with “reasonable” parameter ranges for the optimization in the iterative refinement procedure described above.



...and there's one more problem: cost vs. accuracy trade-off



Accurate methods (usually QM- or DFT-based) are computationally expensive to compute

Force fields are simple to compute, but their accuracy is low compared to coupled-cluster (or even DFT) methods.

The second problem is: how can we obtain accurate, yet fast potentials?

Figure adapted from:

Z. Qiao et al. *J. Chem. Phys.* **153**, 124111 (2020)

(Accuracy computed by Qiao et al. with respect to the Hutchinson conformer benchmark)

Enter machine learning (ML)

The use of ML has become a trend to address issues of automation, pattern recognition, and cost-accuracy trade-off

NUMBER of AI PUBLICATIONS by FIELD of STUDY (EXCLUDING OTHER AI), 2010–21

Source: Center for Security and Emerging Technology, 2021 | Chart: 2022 AI Index Report


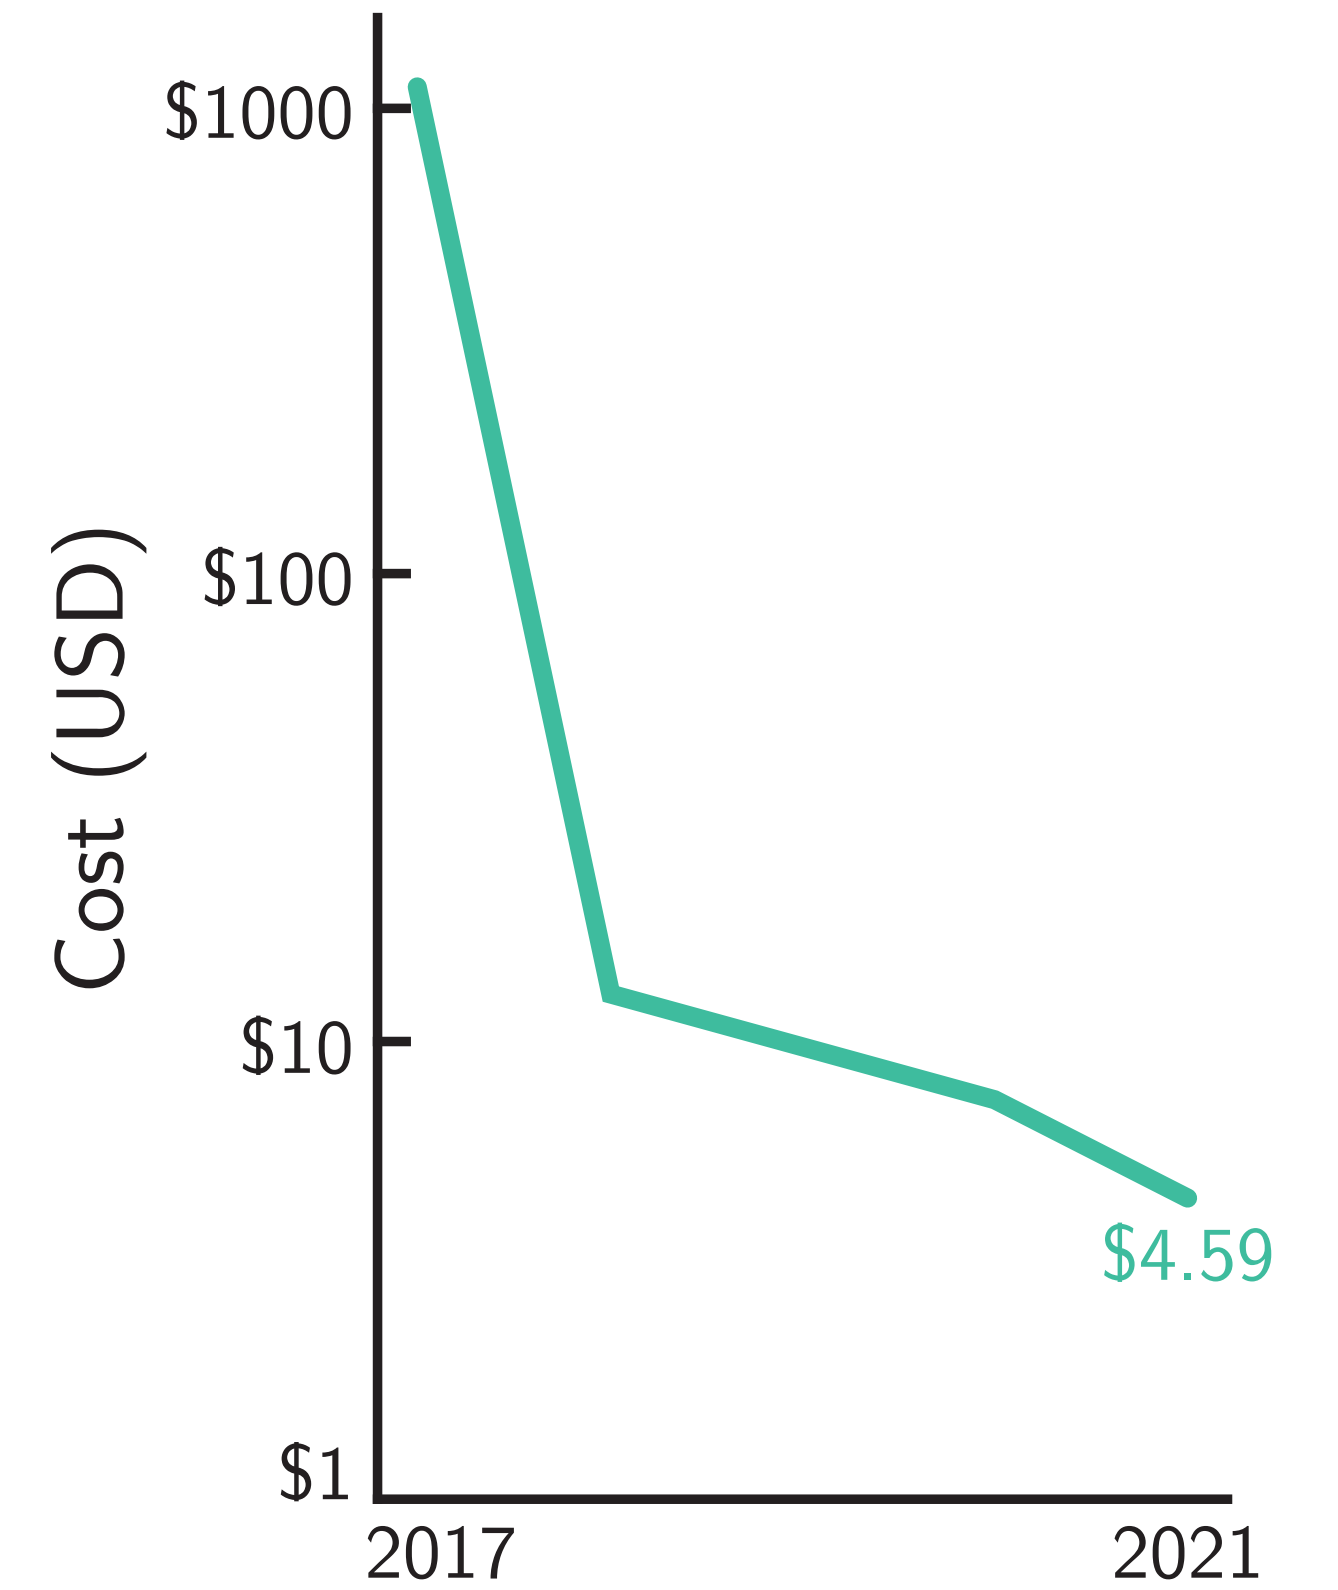


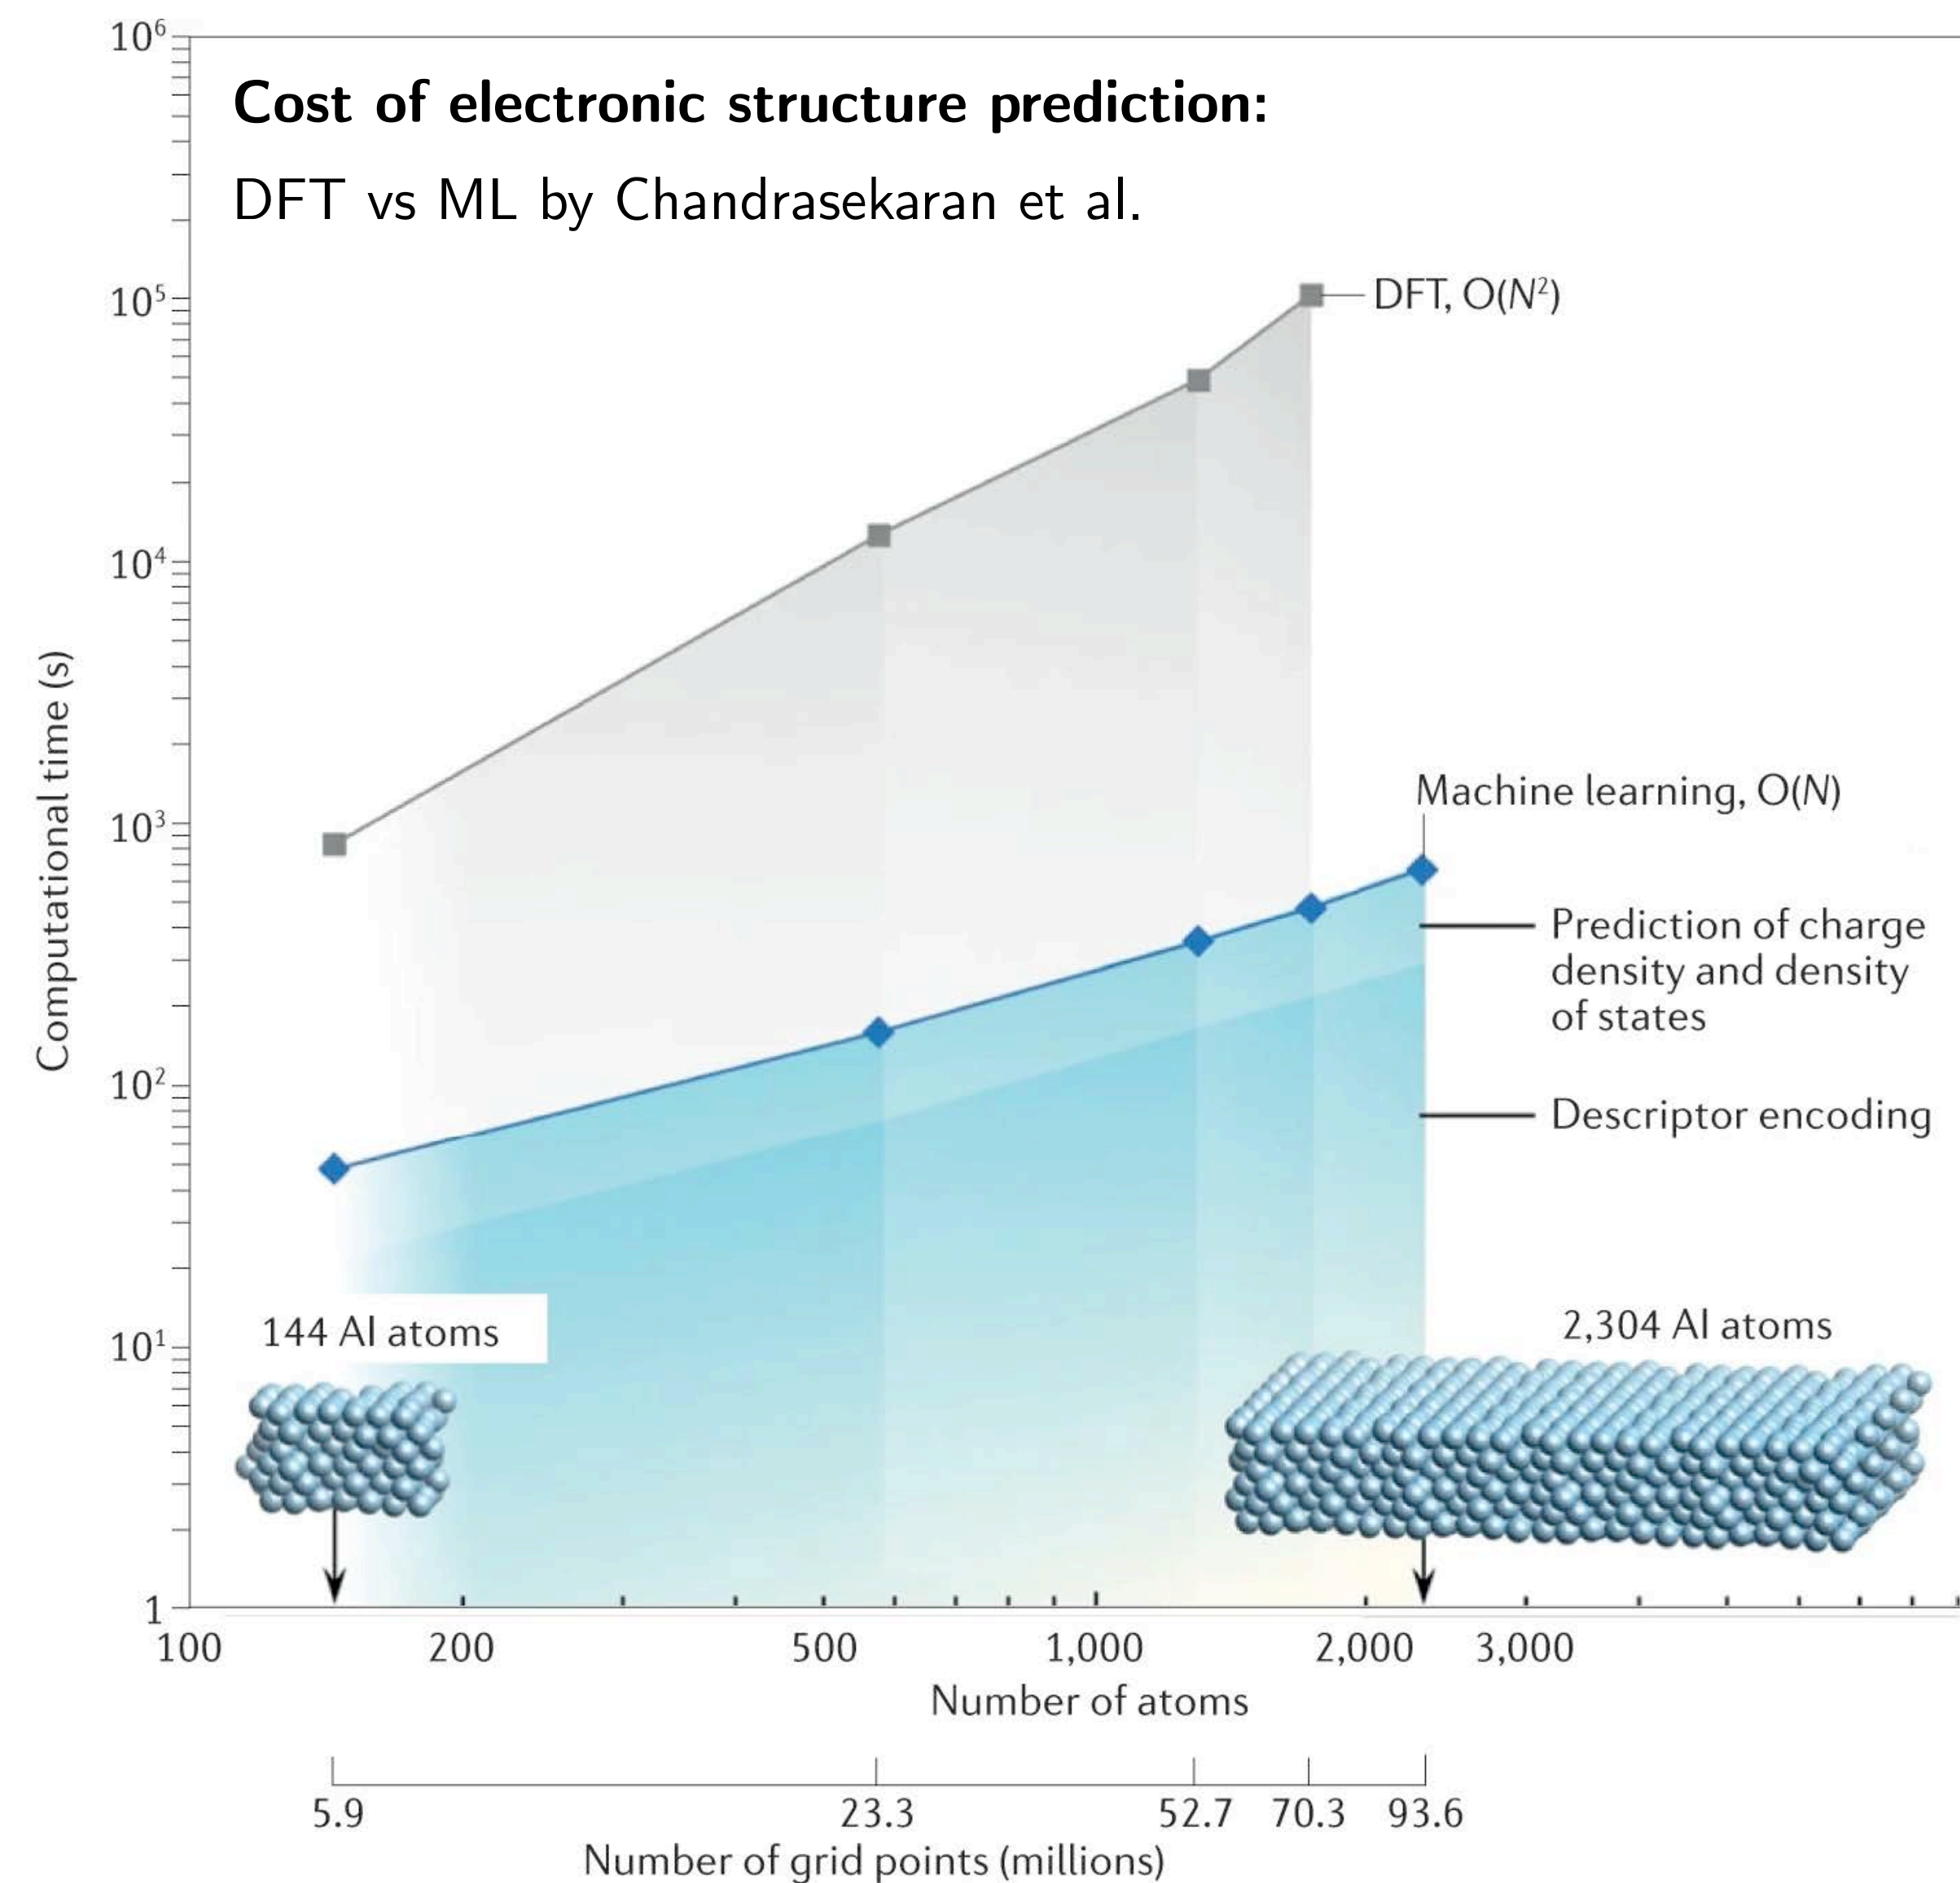
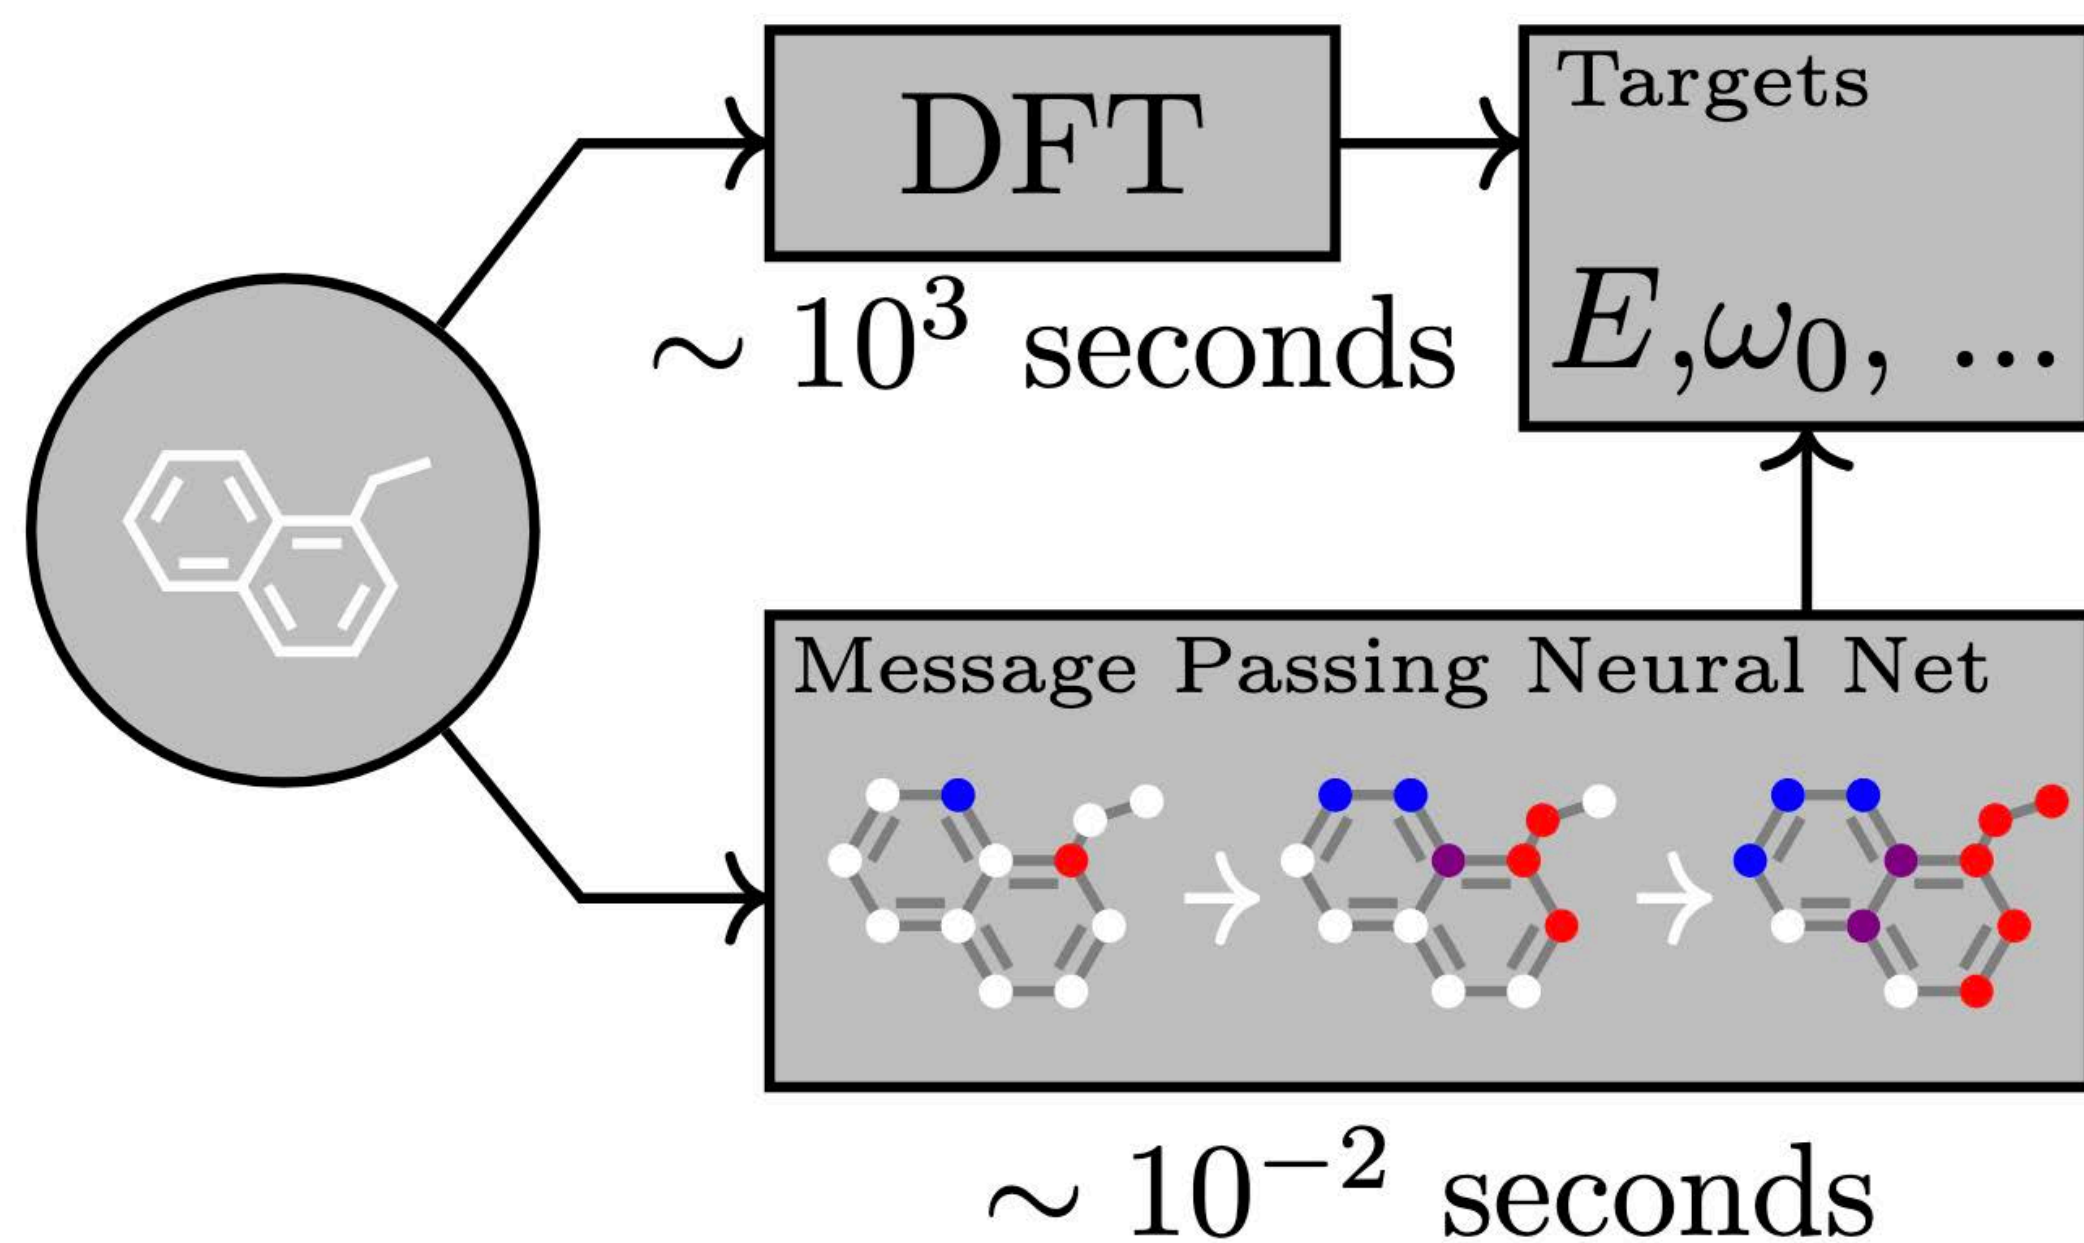
Figure 1.1.3

IMAGENET: training cost (93% acc.)

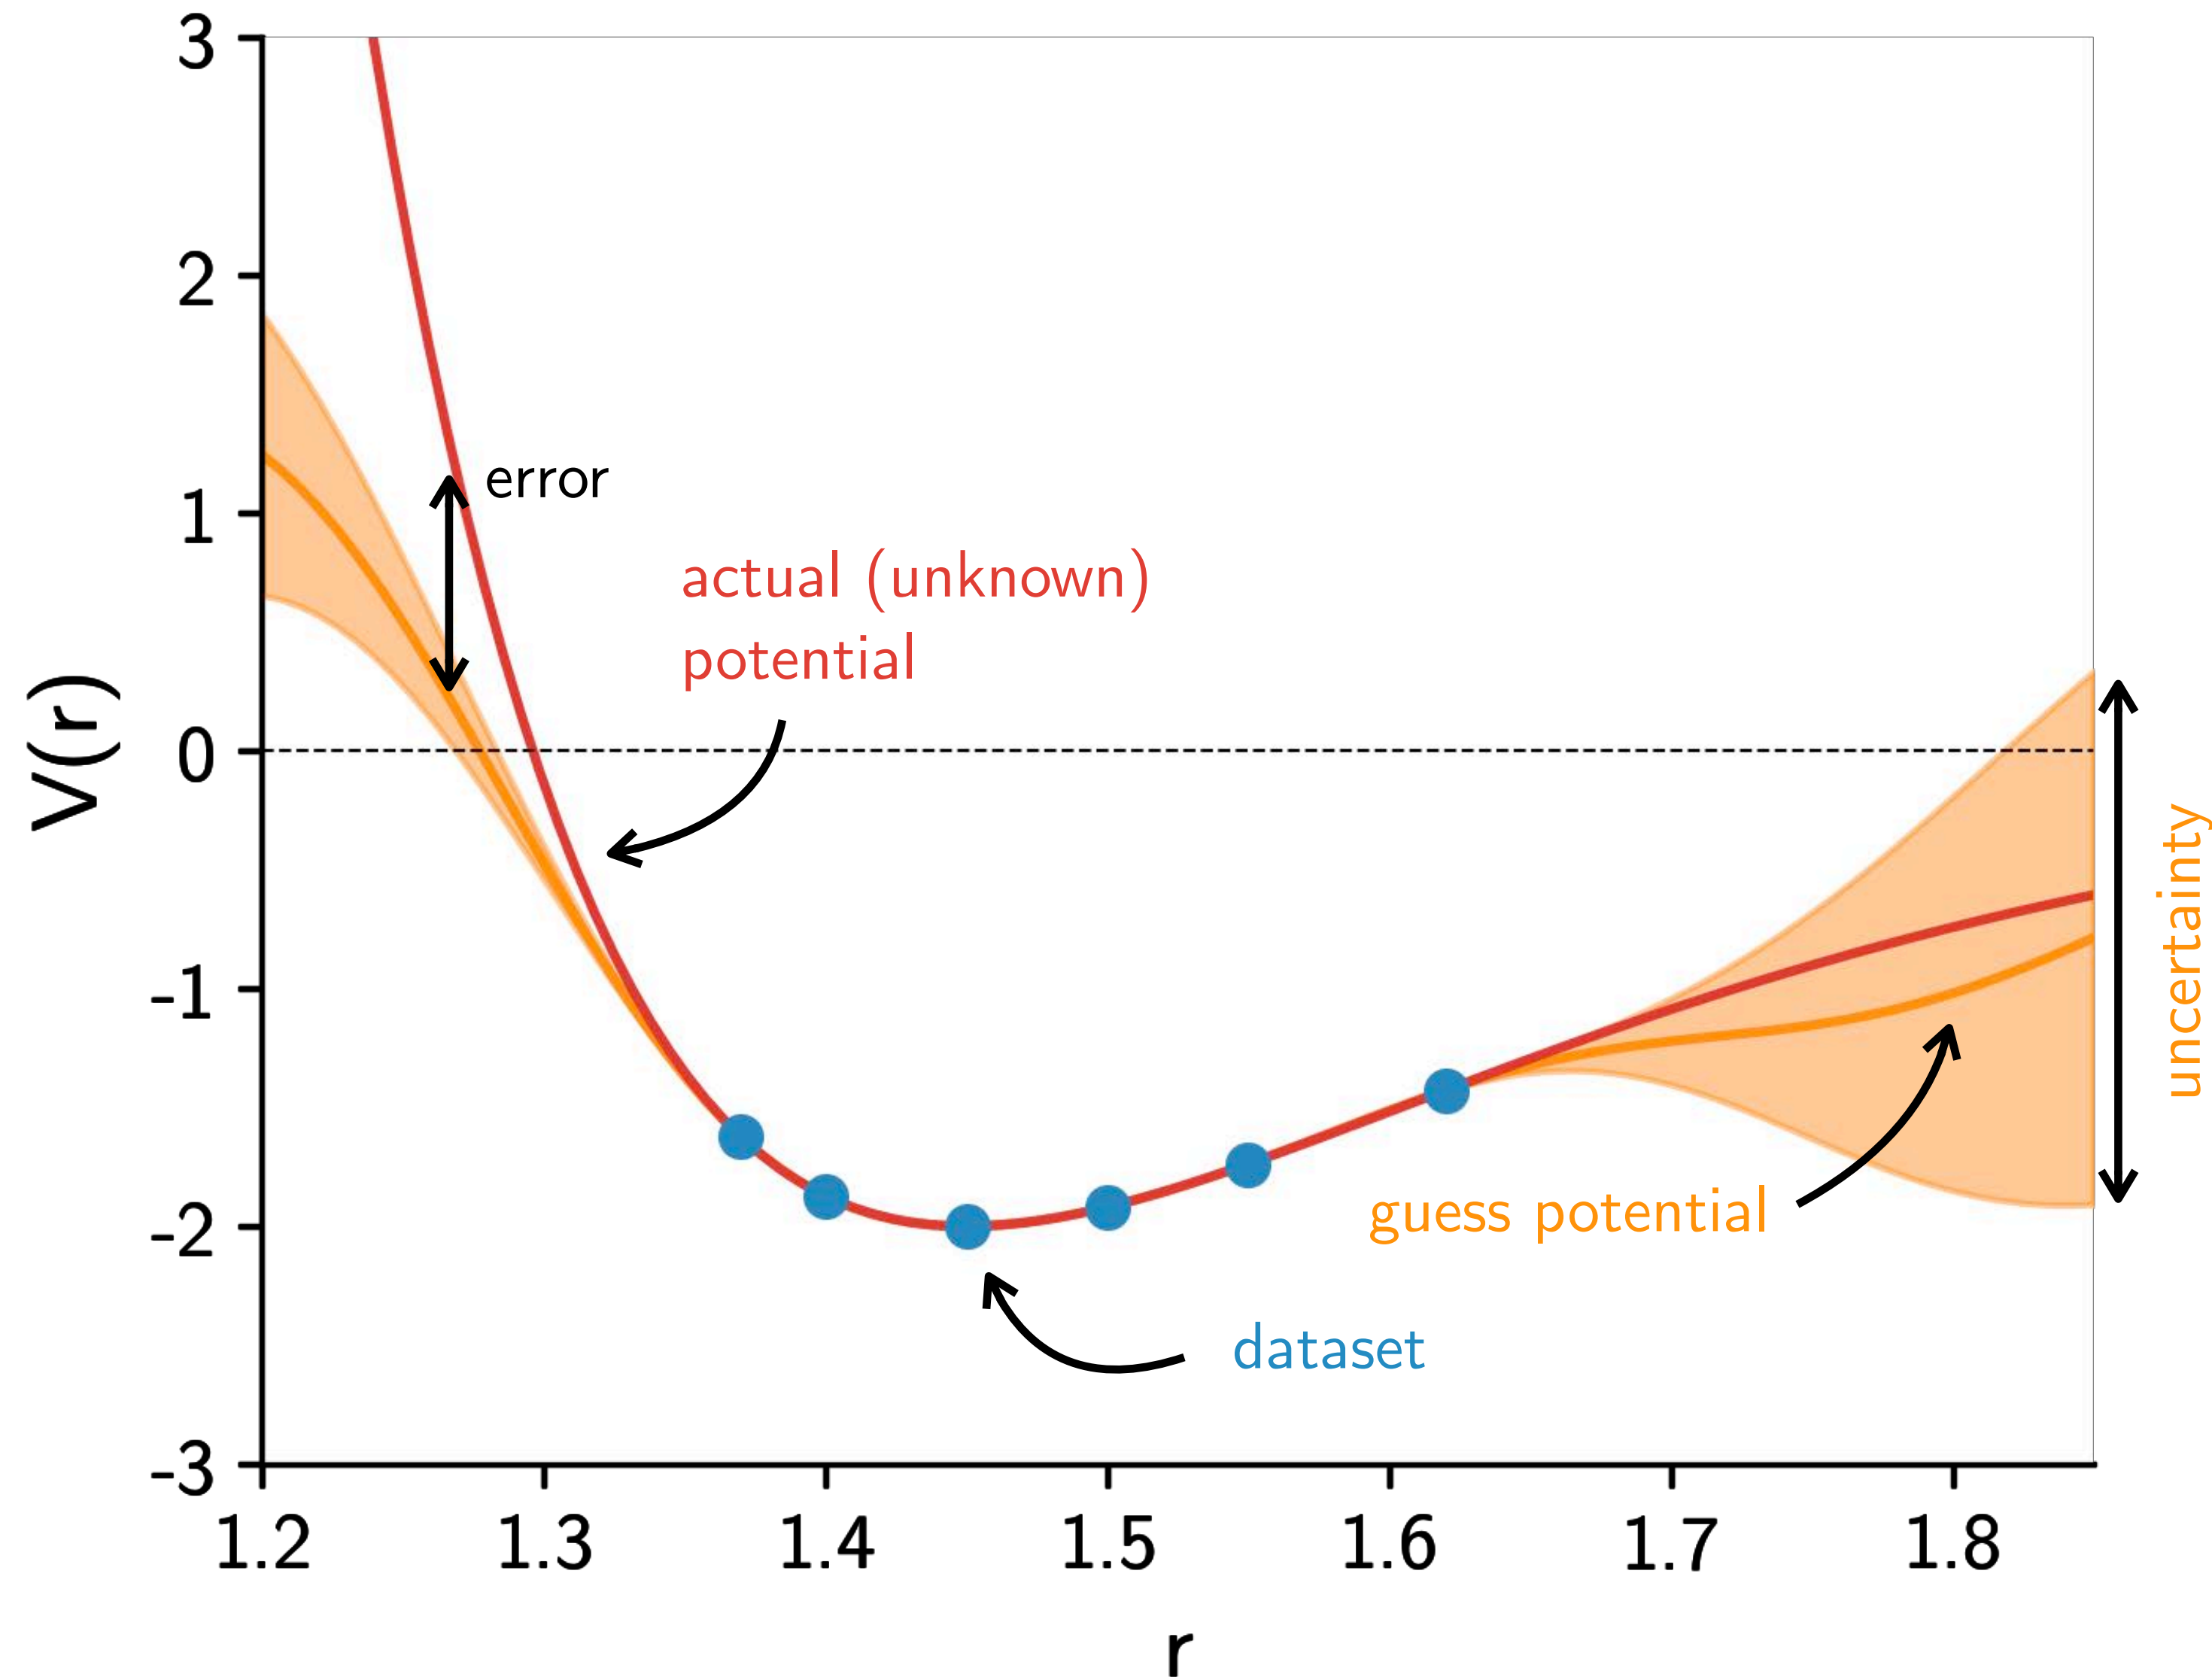


D. Zhang et al. "The AI Index 2022 Annual Report," AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022.

ML has been helping reduce the cost of calculations for materials and chemical systems



Particularly in interatomic potentials, ML helps in fitting to datasets

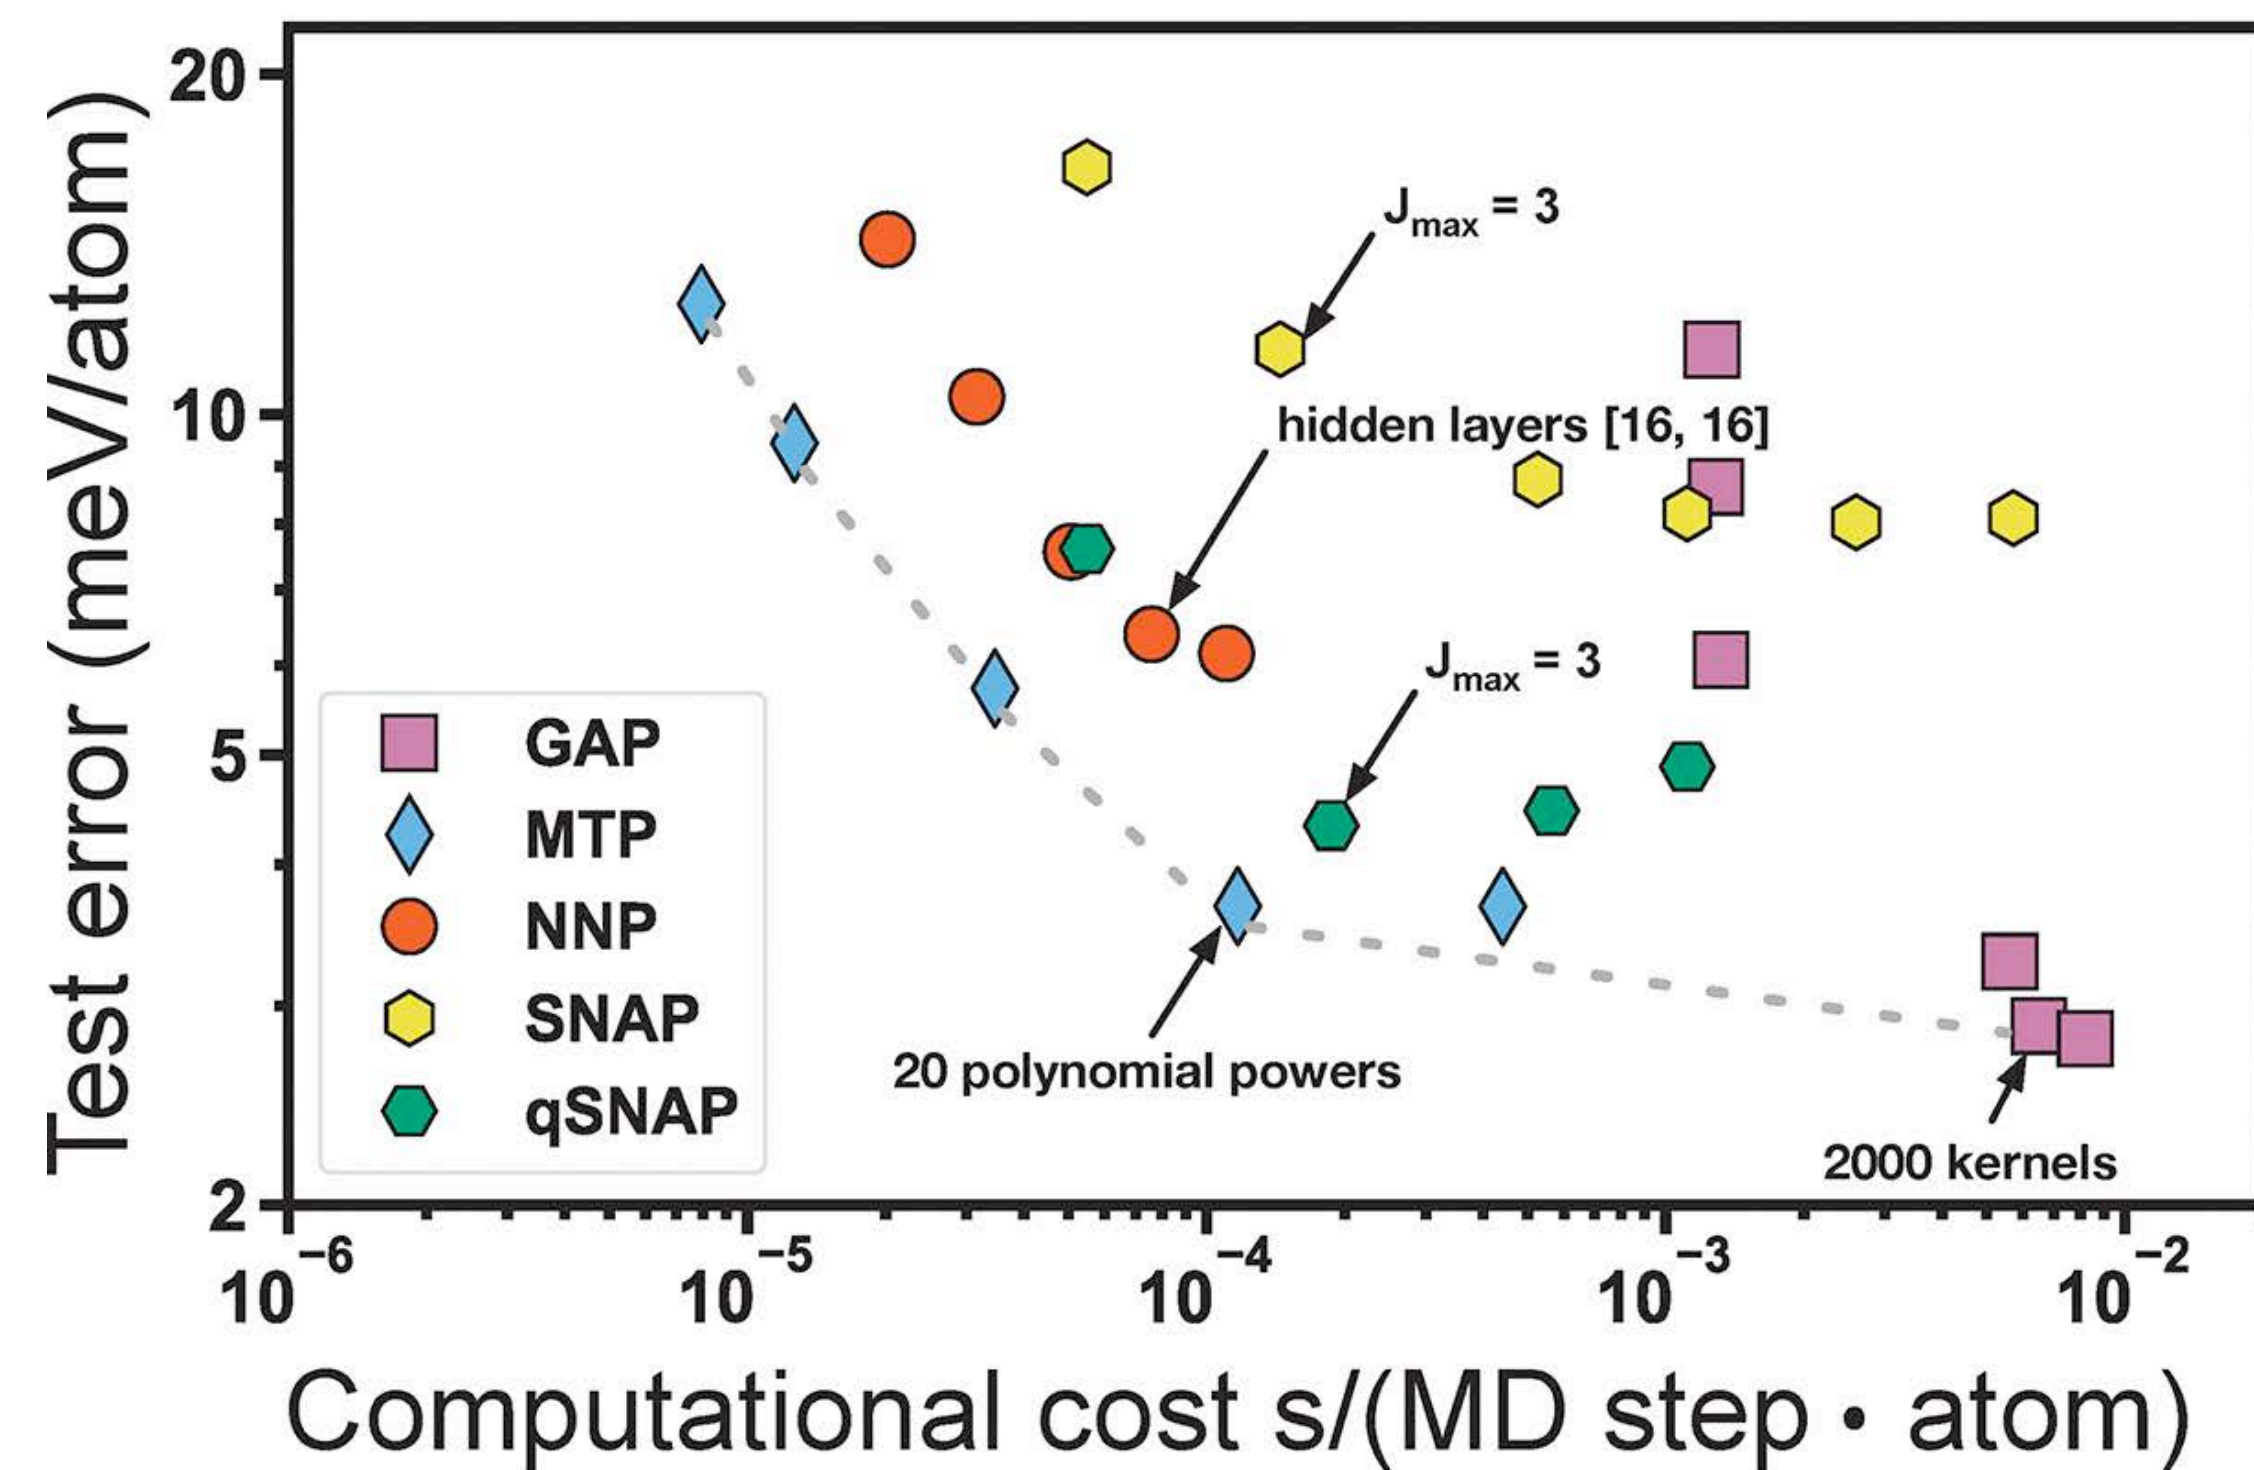
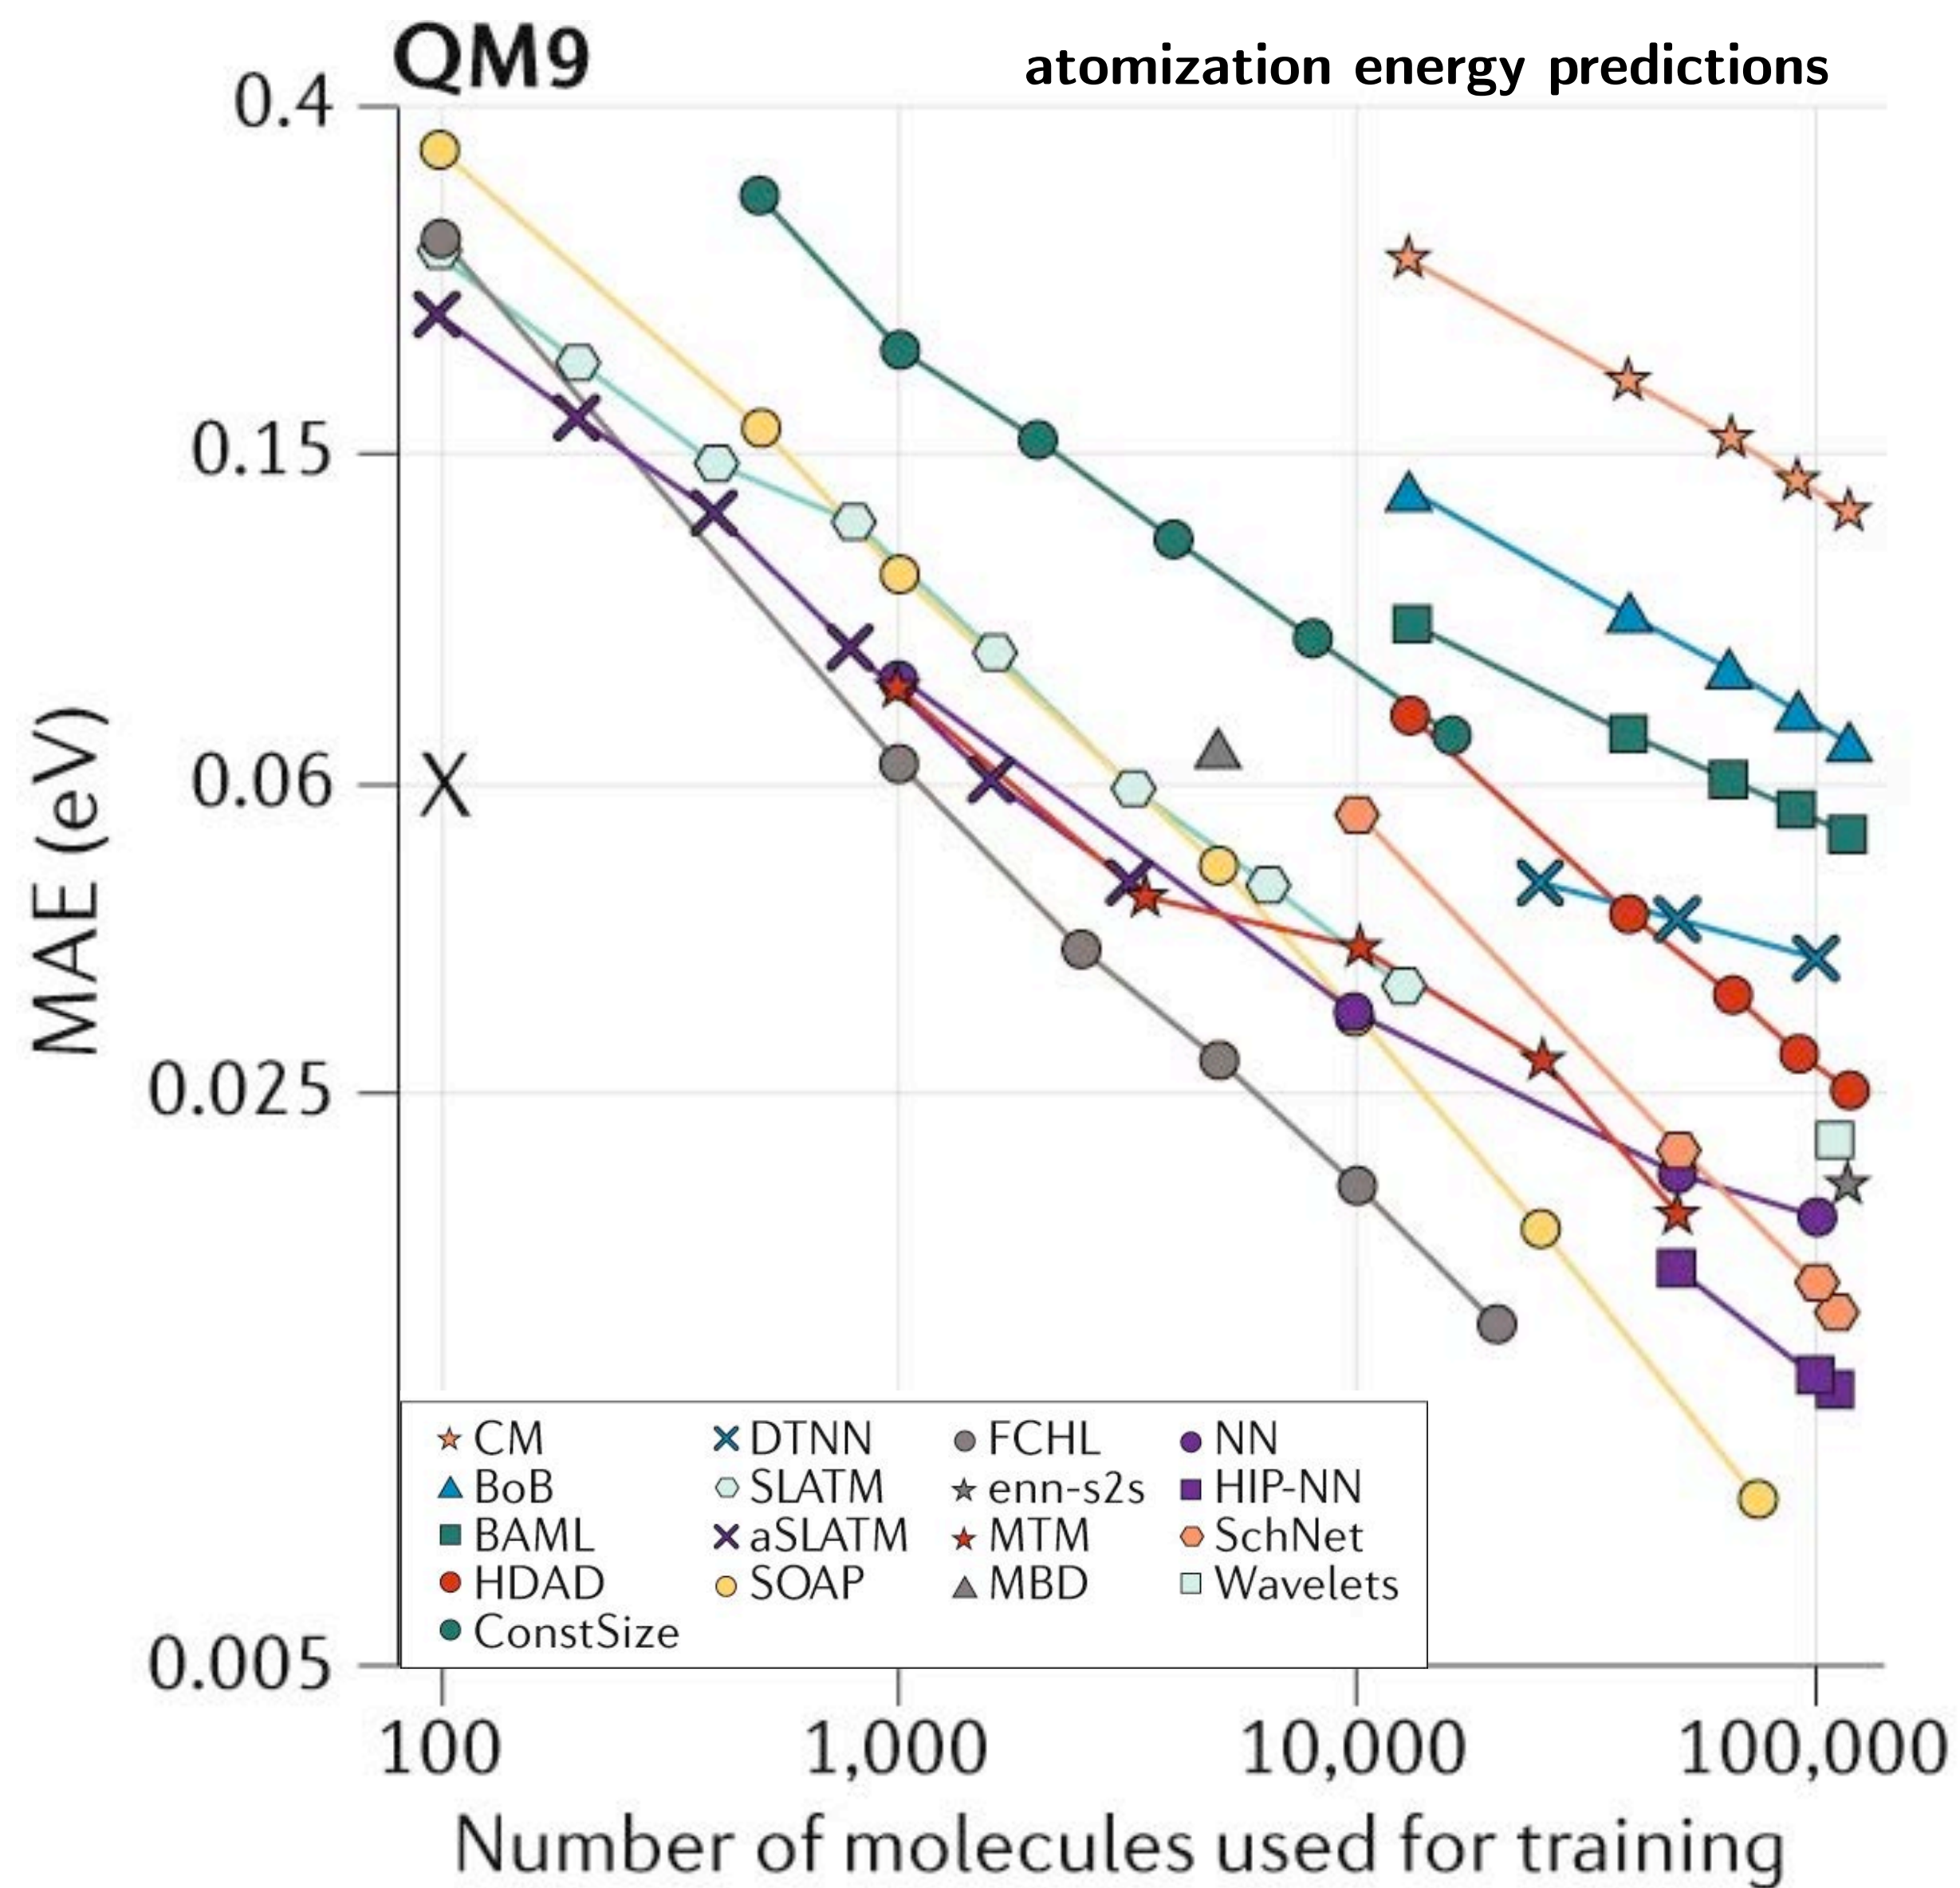


ML methods enable a fitting to guess potentials based on a given dataset of interest.

When implemented, this approach can automate the process of finding functional forms that fit to the data and bypasses functional forms that can be less accurate.


But which ML methods should we use to implement force fields?

There are many ML methods and implementations possible...



...but the most popular ones are:

Linear Methods



Polynomial on many-body terms




Simple and fast



Relies crafting a representation for the inputs

Kernel/Gaussian Process Regression



Computes an explicit similarity between points

Fewer data points
(+uncertainty for GPR)

$O(N^3)$ complexity for training for GPR

Neural Networks (NNs)



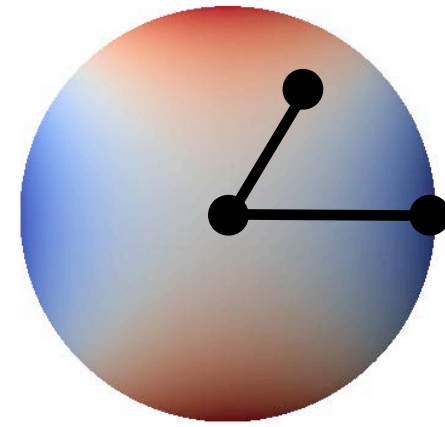
“Universal approximator” with non-linear mappings

High accuracy

Large number of trainable parameters

Some examples in the literature:

Linear Methods




SNAP (Thompson et al.)

MTP (Shapeev)

ACE (Drautz, Kovács et al.)

Kernel/Gaussian Process Regression



GPR:

GAP (Bartok et al.)

MLOTF (Li et al.)

FLARE (Vandermause et al.)

Other kernels with:


sGDML (Chmiela et al.)

FCHL repres. (Faber et al.)

Coulomb matrices (Rupp et al.)

this lecture:

Neural Networks (NNs)



Behler-Parrinello

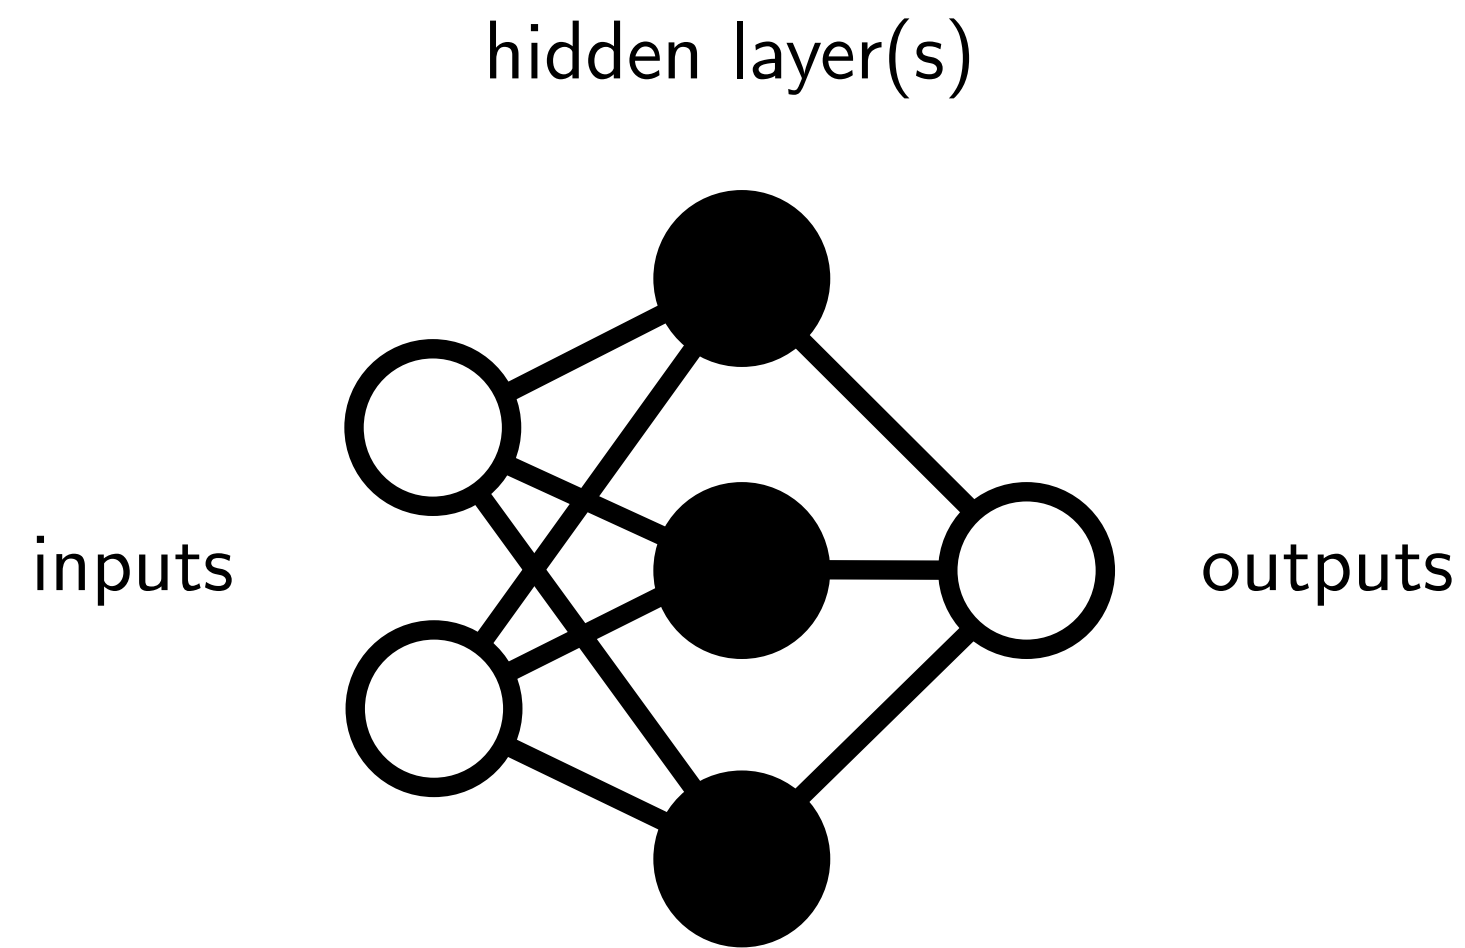
Representation + NN
(DeepMD, ANI etc.)

Deep learning-based NNFF
(SchNet etc.)

Deep learning + equivariance
(NequIP, PaiNN etc.)

Deep learning + many-body
expansion (MACE etc.)

One-slide neural network refresher



$$f(\mathbf{X}) = \sigma(W\mathbf{X} + \mathbf{b})$$

↑ input ↑ weight matrix ↑ bias
 ↖ non-linearity

$$\hat{y} = f(\mathbf{X})$$

↑
estimated output

Training of neural networks requires setting a loss function (e.g., for regressors, no regularization):

$$\mathcal{L} = \mathbb{E}_{\mathbf{X} \sim P(\mathbf{X})} [\|\hat{y} - y\|^2]$$

which updates the weights using the backpropagation algorithm and gradient descent:

$$w_{ij}^{(n+1)} = w_{ij}^{(n)} - \alpha \frac{\partial \mathcal{L}}{\partial w_{ij}}$$

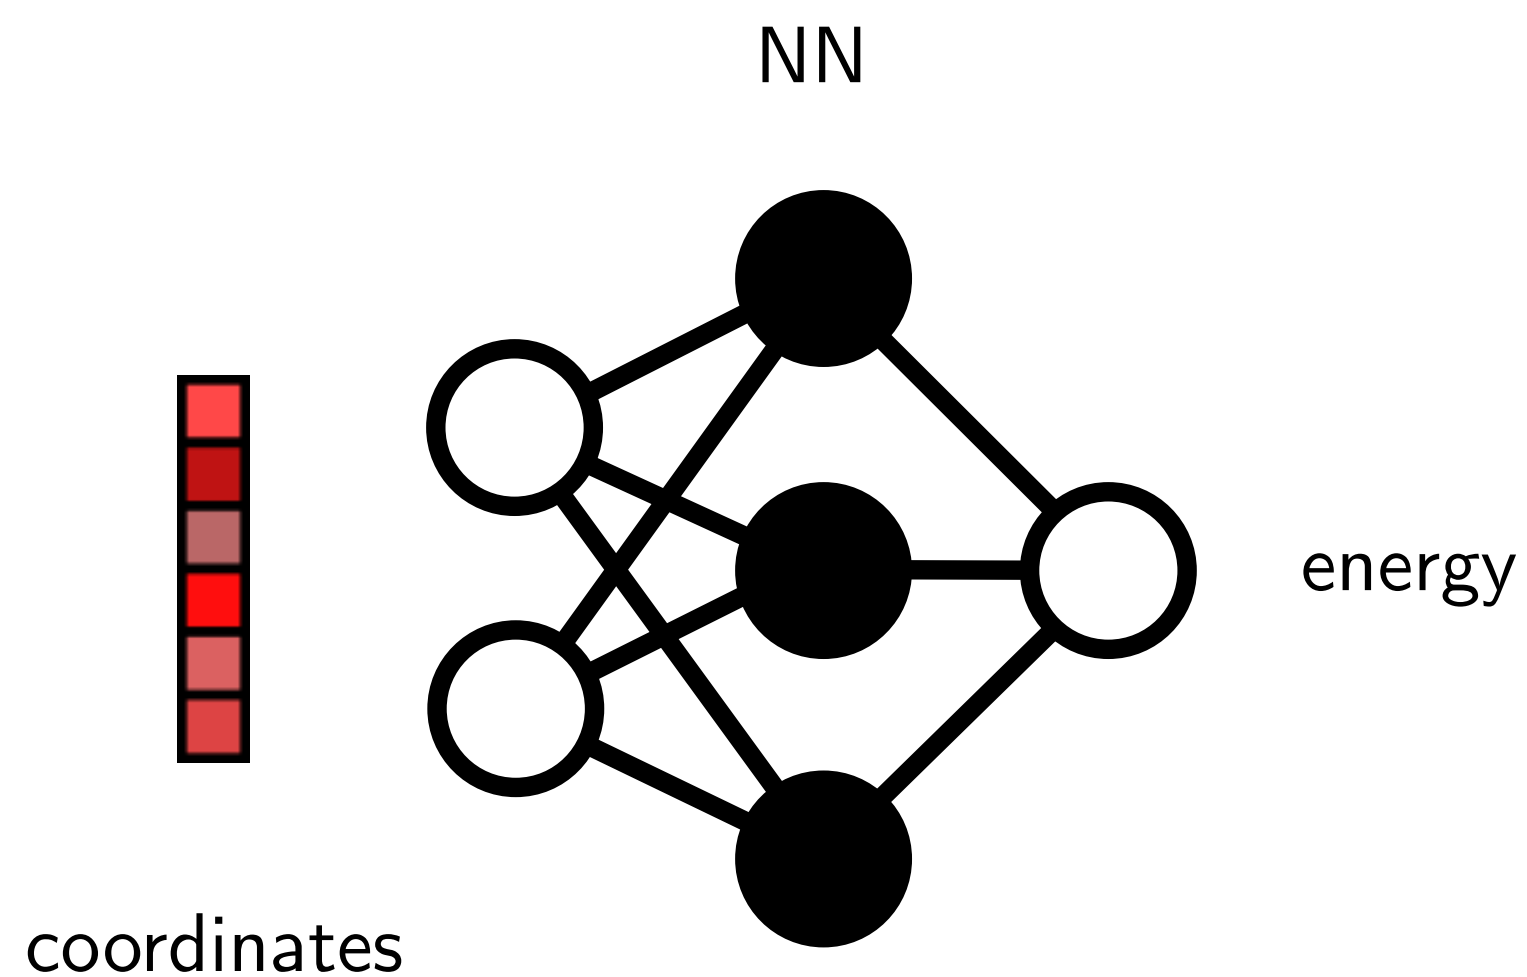
↑ weight at iteration (n + 1) ↑ learning rate ↖ loss gradient

Mini Tutorial: NNs are “universal approximators”

2. NN Interatomic Potentials

Fitting PESes with neural networks

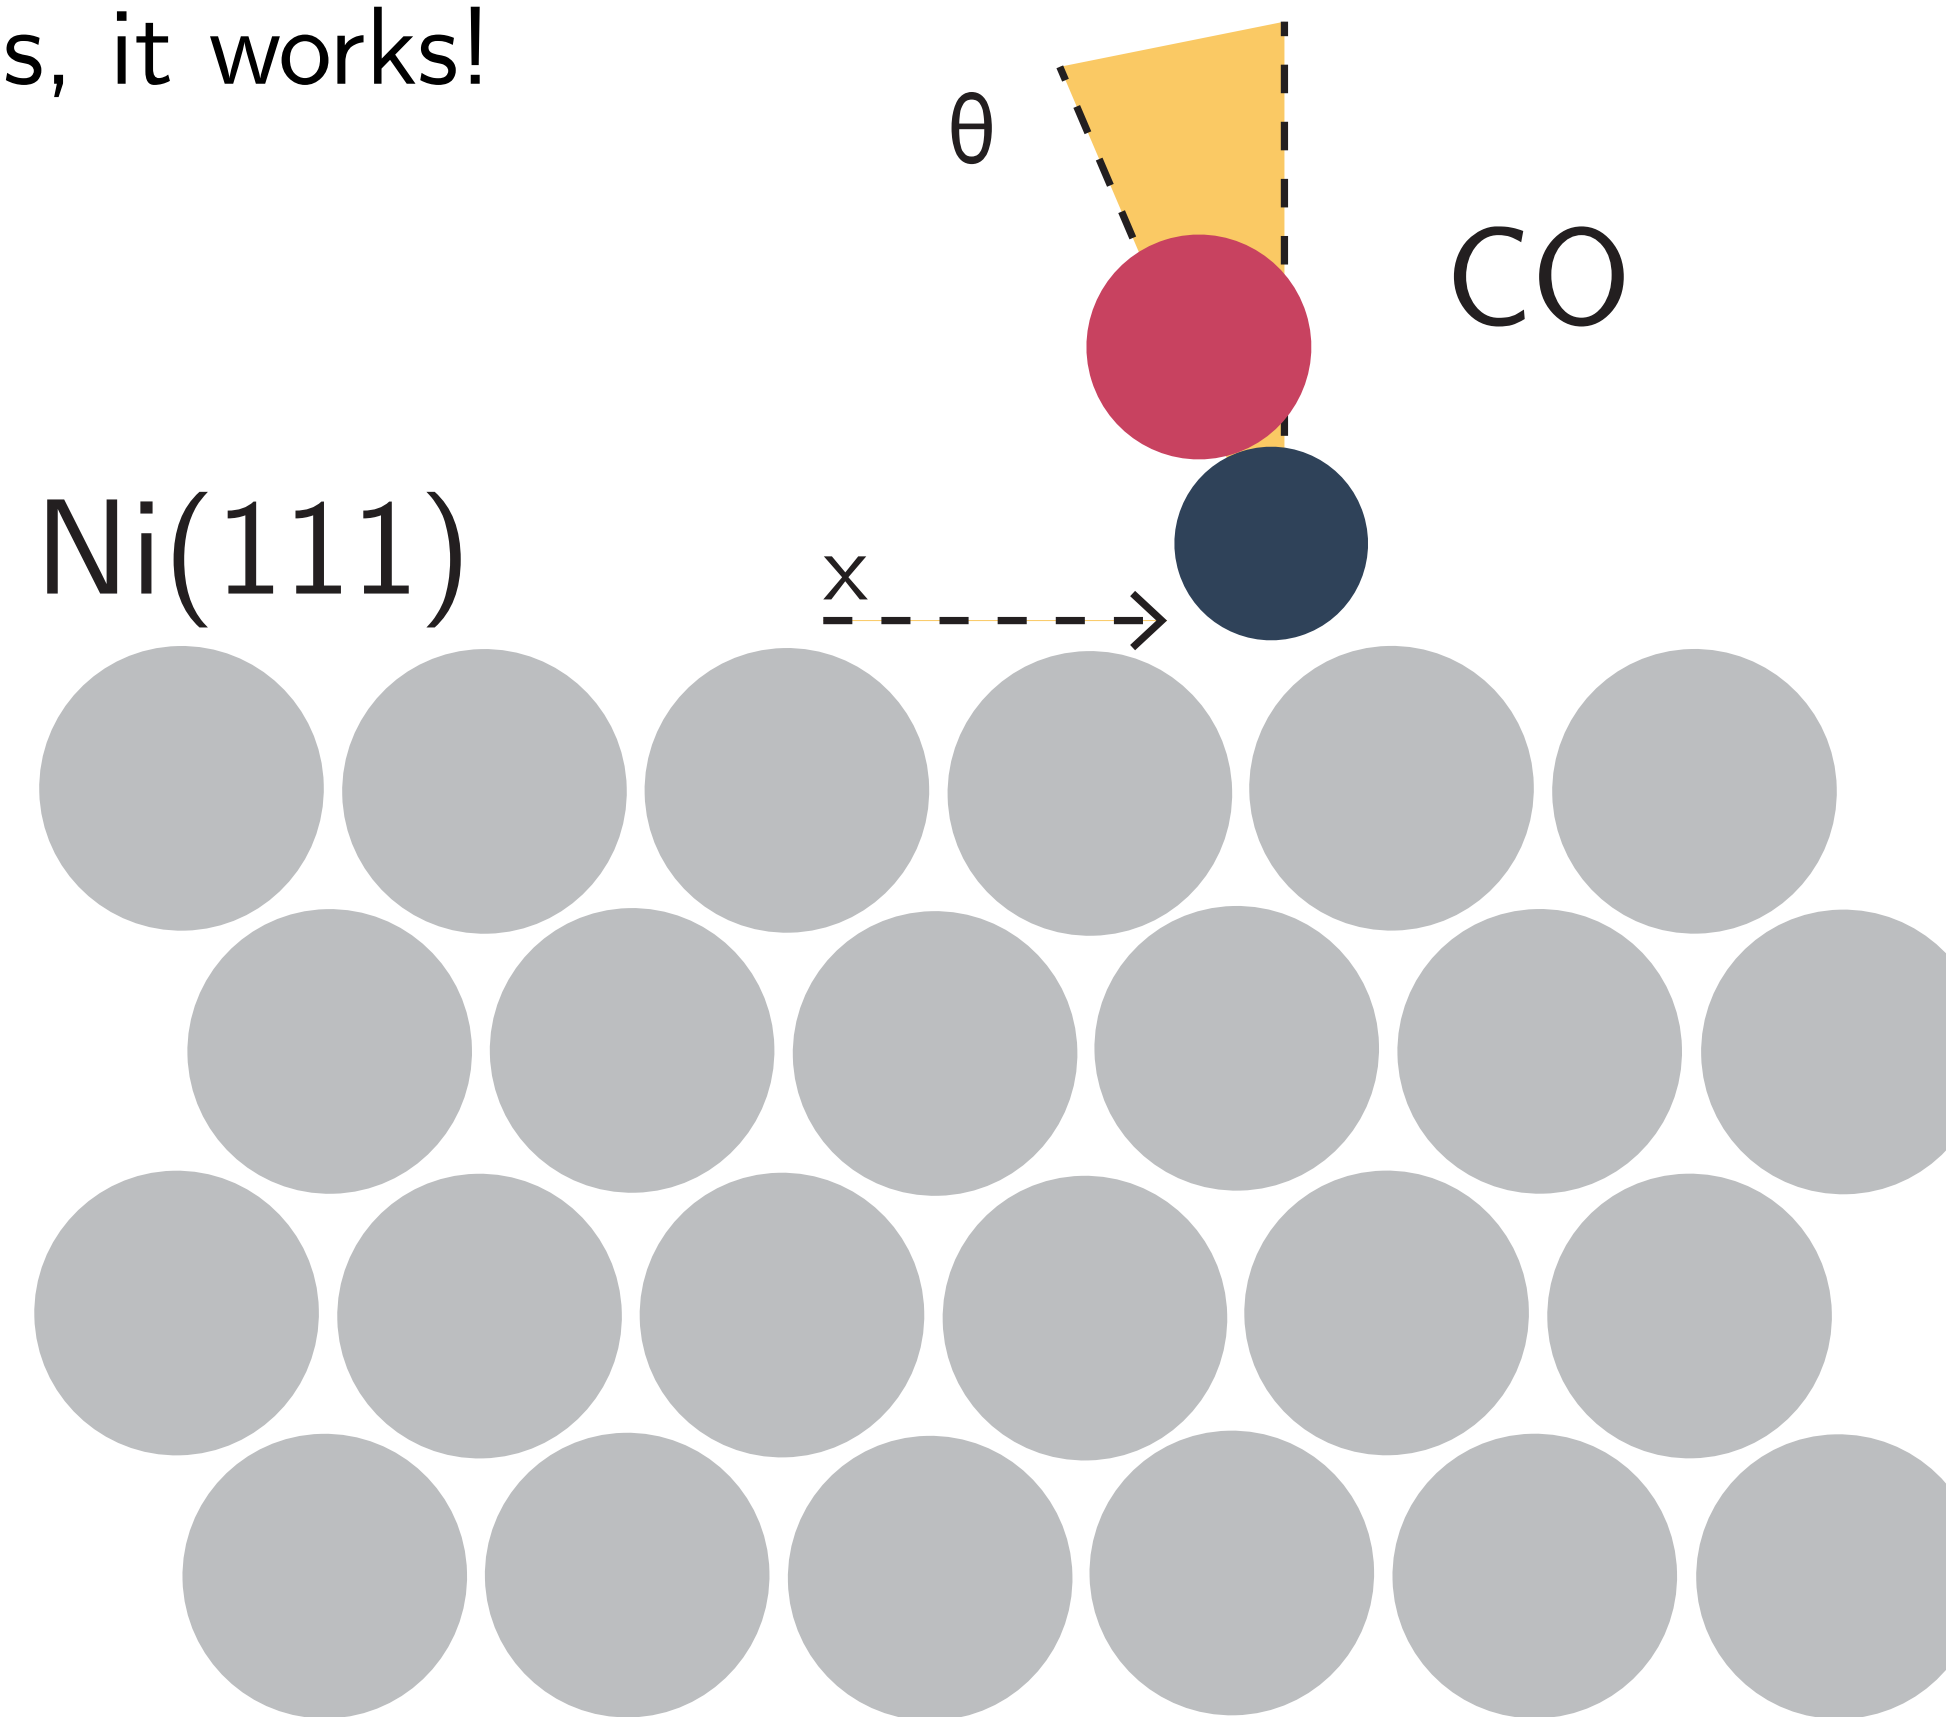
Simplest possible approach in fitting a PES using a neural network



The NN is trained to map input coordinates to energy for a given dataset.

Does this work?

Yes, it works!



T. Blank et al. *J. Chem. Phys.* **103**, 4129 (1995)

What are the advantages/
disadvantages of this approach?

What are the advantages/disadvantages?

 **Advantages**

Fast

No need for extra
information

Easy interpretation


 **Disadvantages**

No invariance to
translation/rotation

(and/or) No invariance to
permutation


No forces

Does not scale to larger
systems



The neighborhood of each atom is encoded into symmetry functions, which embed the rototranslational invariance of the system.

Atom-centered
symmetry
functions



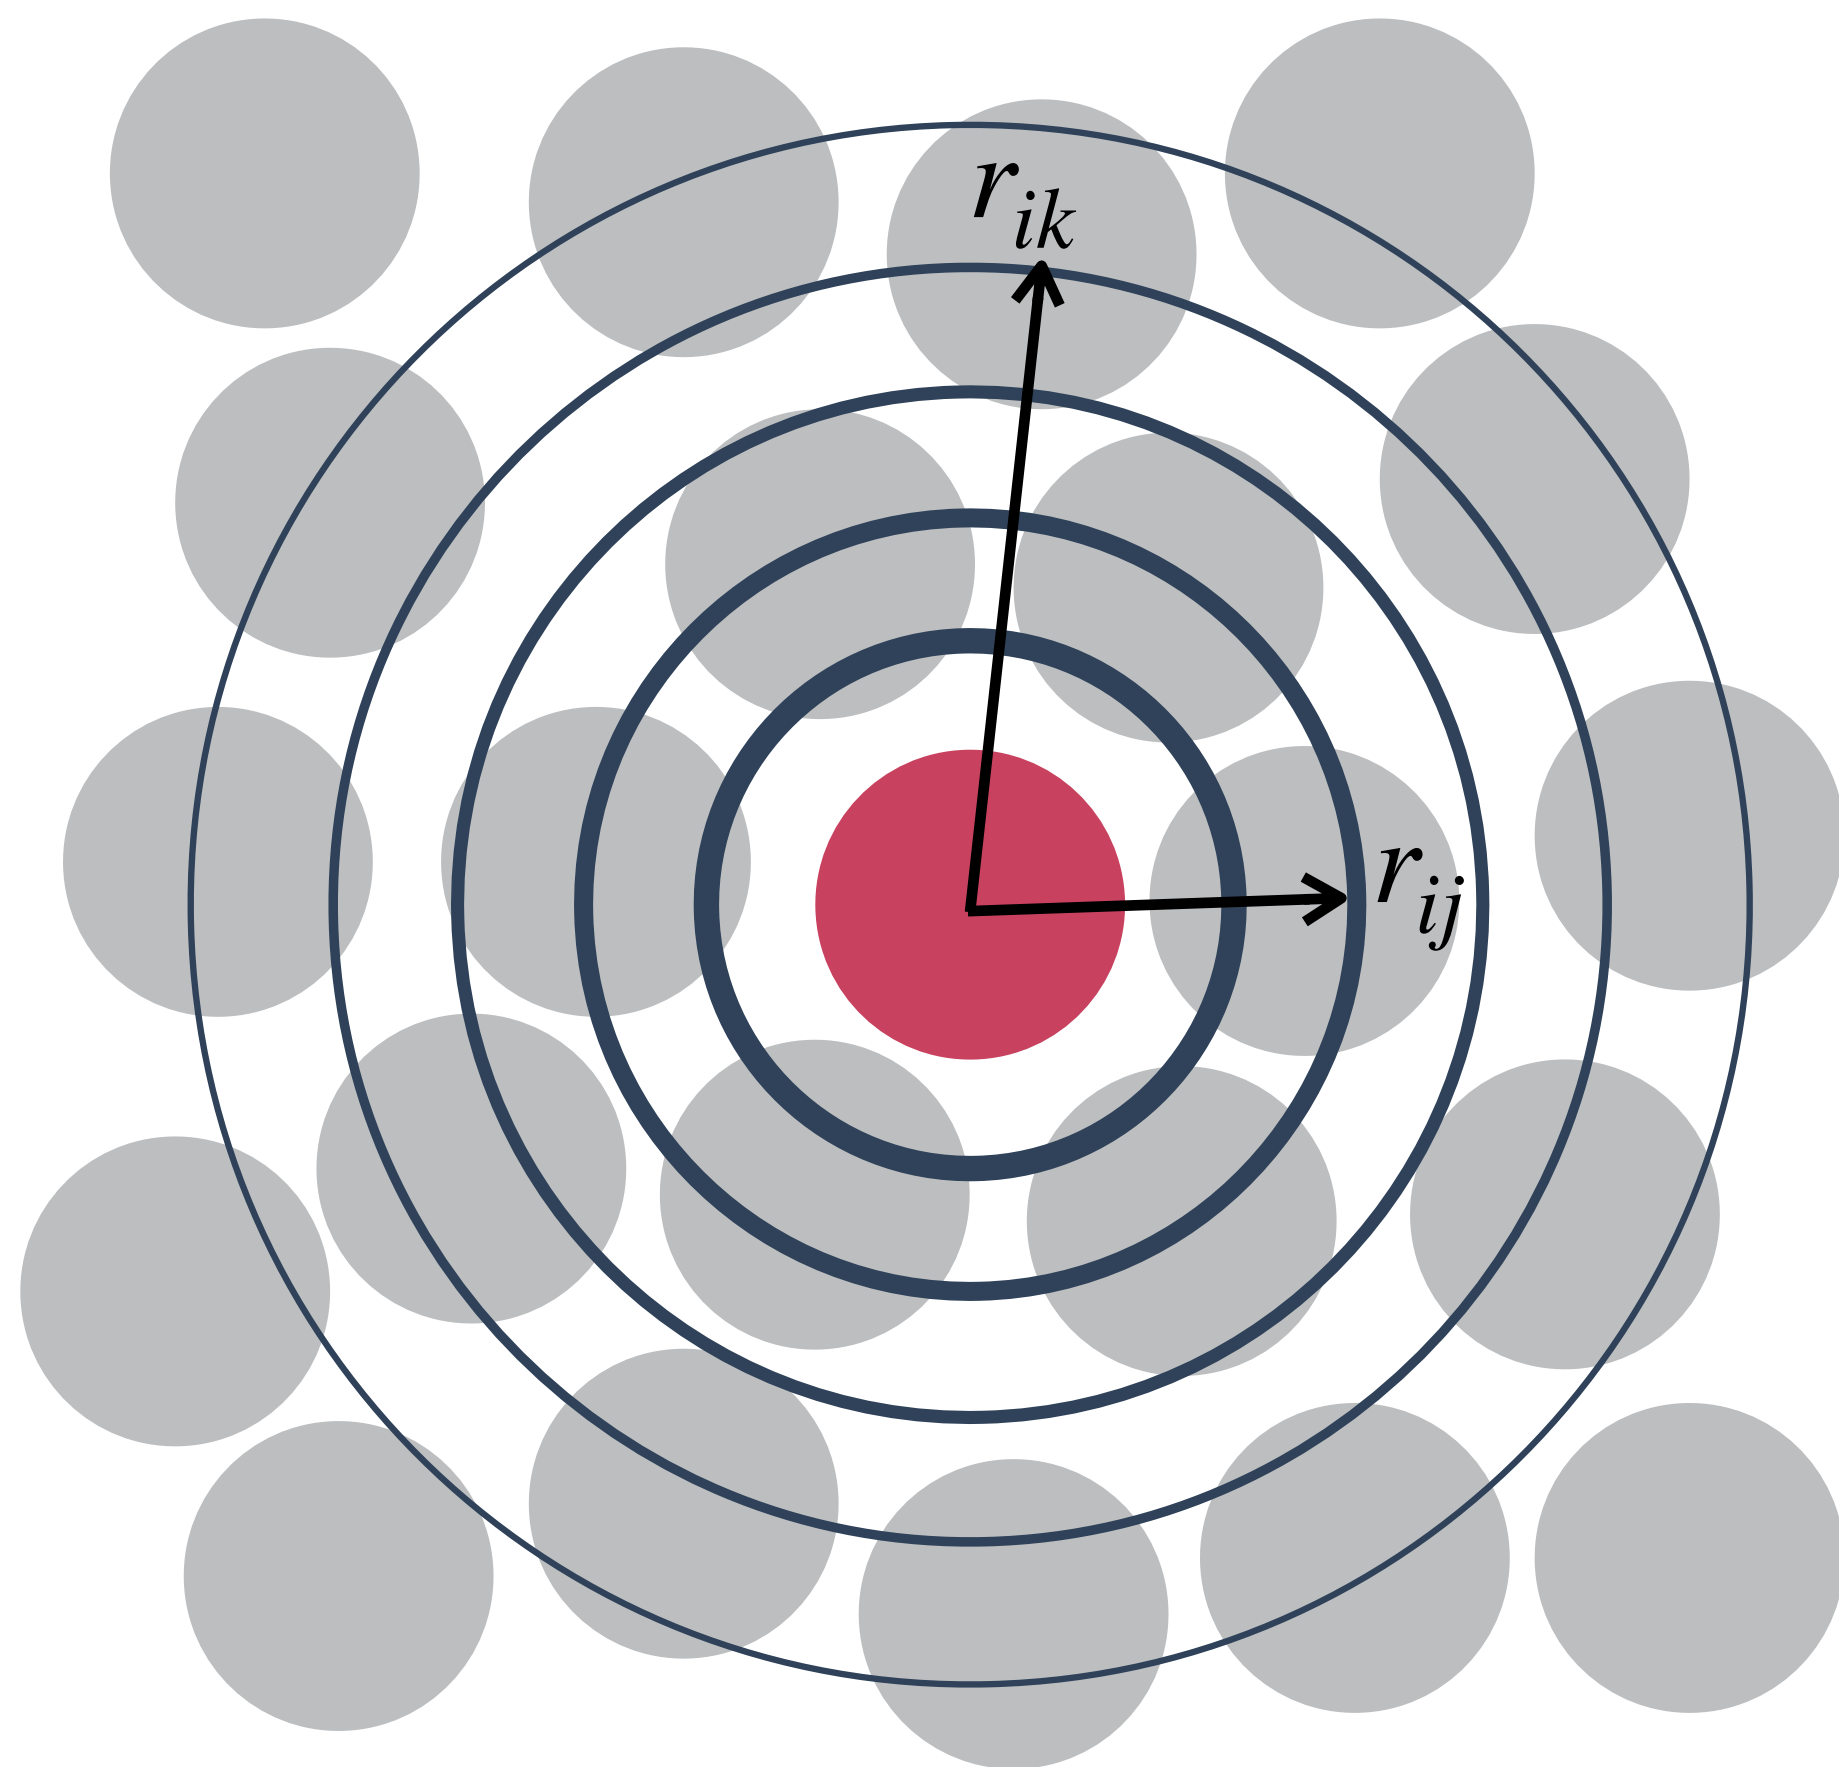
Per-atom energy
contribution

The final energy is computed as:

$$E = \sum_i E_i$$

Think about this: does dividing the energy make sense?

But how to define symmetry functions, now?



For each atom, N-body symmetry functions are calculated with different hyperparameters:

$$(\eta, r_s, \zeta)$$

Symmetry functions

cutoff

$$f_c(r_{ij}) = \frac{1}{2} \cos\left(\frac{\pi r_{ij}}{r_c}\right) + \frac{1}{2}, \quad r_{ij} < r_c$$

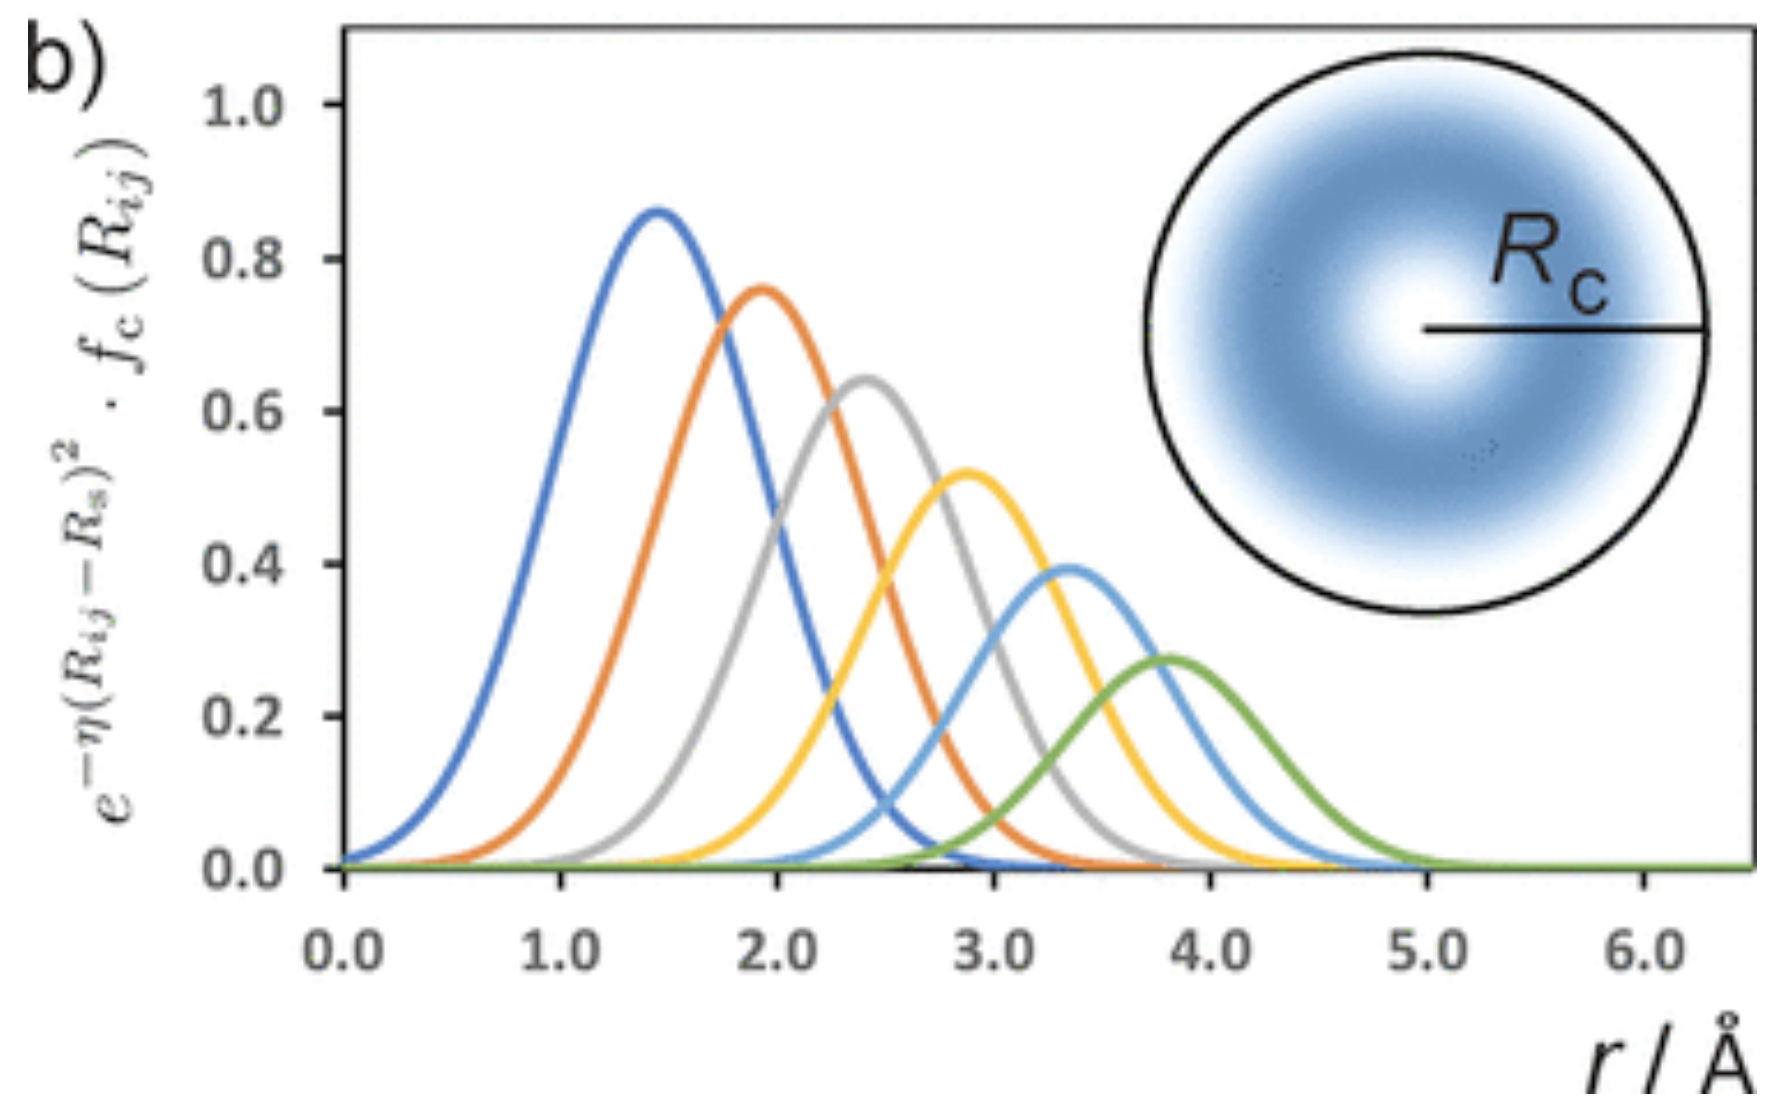

radial

$$G_i^1 = \sum_{i \neq j} e^{-\eta(r_{ij}-r_s)^2} f_c(r_{ij})$$

angular

$$G_i^2 = 2^{1-\zeta} \sum_{i \neq j, k} (1 + \lambda \cos \theta_{ijk})^\zeta \\ \times e^{-\eta(r_{ij}^2 + r_{ik}^2 + r_{jk}^2)} \\ \times f_c(r_{ij}) f_c(r_{ik}) f_c(r_{jk})$$

Visualizing these symmetry functions



Symmetry functions

cutoff

$$f_c(r_{ij}) = \frac{1}{2} \cos\left(\frac{\pi r_{ij}}{r_c}\right) + \frac{1}{2}, \quad r_{ij} < r_c$$



radial

$$G_i^1 = \sum_{i \neq j} e^{-\eta(r_{ij}-r_s)^2} f_c(r_{ij})$$

angular


$$G_i^2 = 2^{1-\zeta} \sum_{i \neq j, k} (1 + \lambda \cos \theta_{ijk})^\zeta \\ \times e^{-\eta(r_{ij}^2 + r_{ik}^2 + r_{jk}^2)} \\ \times f_c(r_{ij}) f_c(r_{ik}) f_c(r_{jk})$$

Results for silicon melt at 3000 K



The NN better approximates the structural properties of Si melt than other potentials


Another benefit of using NN potentials: much better scaling with system size



Because DFT scales with the cube of the number of electrons, it is impractical to perform simulations for very large systems (more than a few hundreds of atoms)

On the other hand, evaluating energies in NNIPs often scales linearly with the number of data points, and can be easily parallelized considering local potentials.

But the locality of the descriptor is also a shortcoming



Because of the locality of the descriptor, the final model cannot fit to interactions such as long-range potentials.

Increasing the cutoff drastically increases the number of atomic environments that have to be sampled.

For multi-element systems, this is even harder.

In the figure above, only atoms within R_c of the central atom are considered.

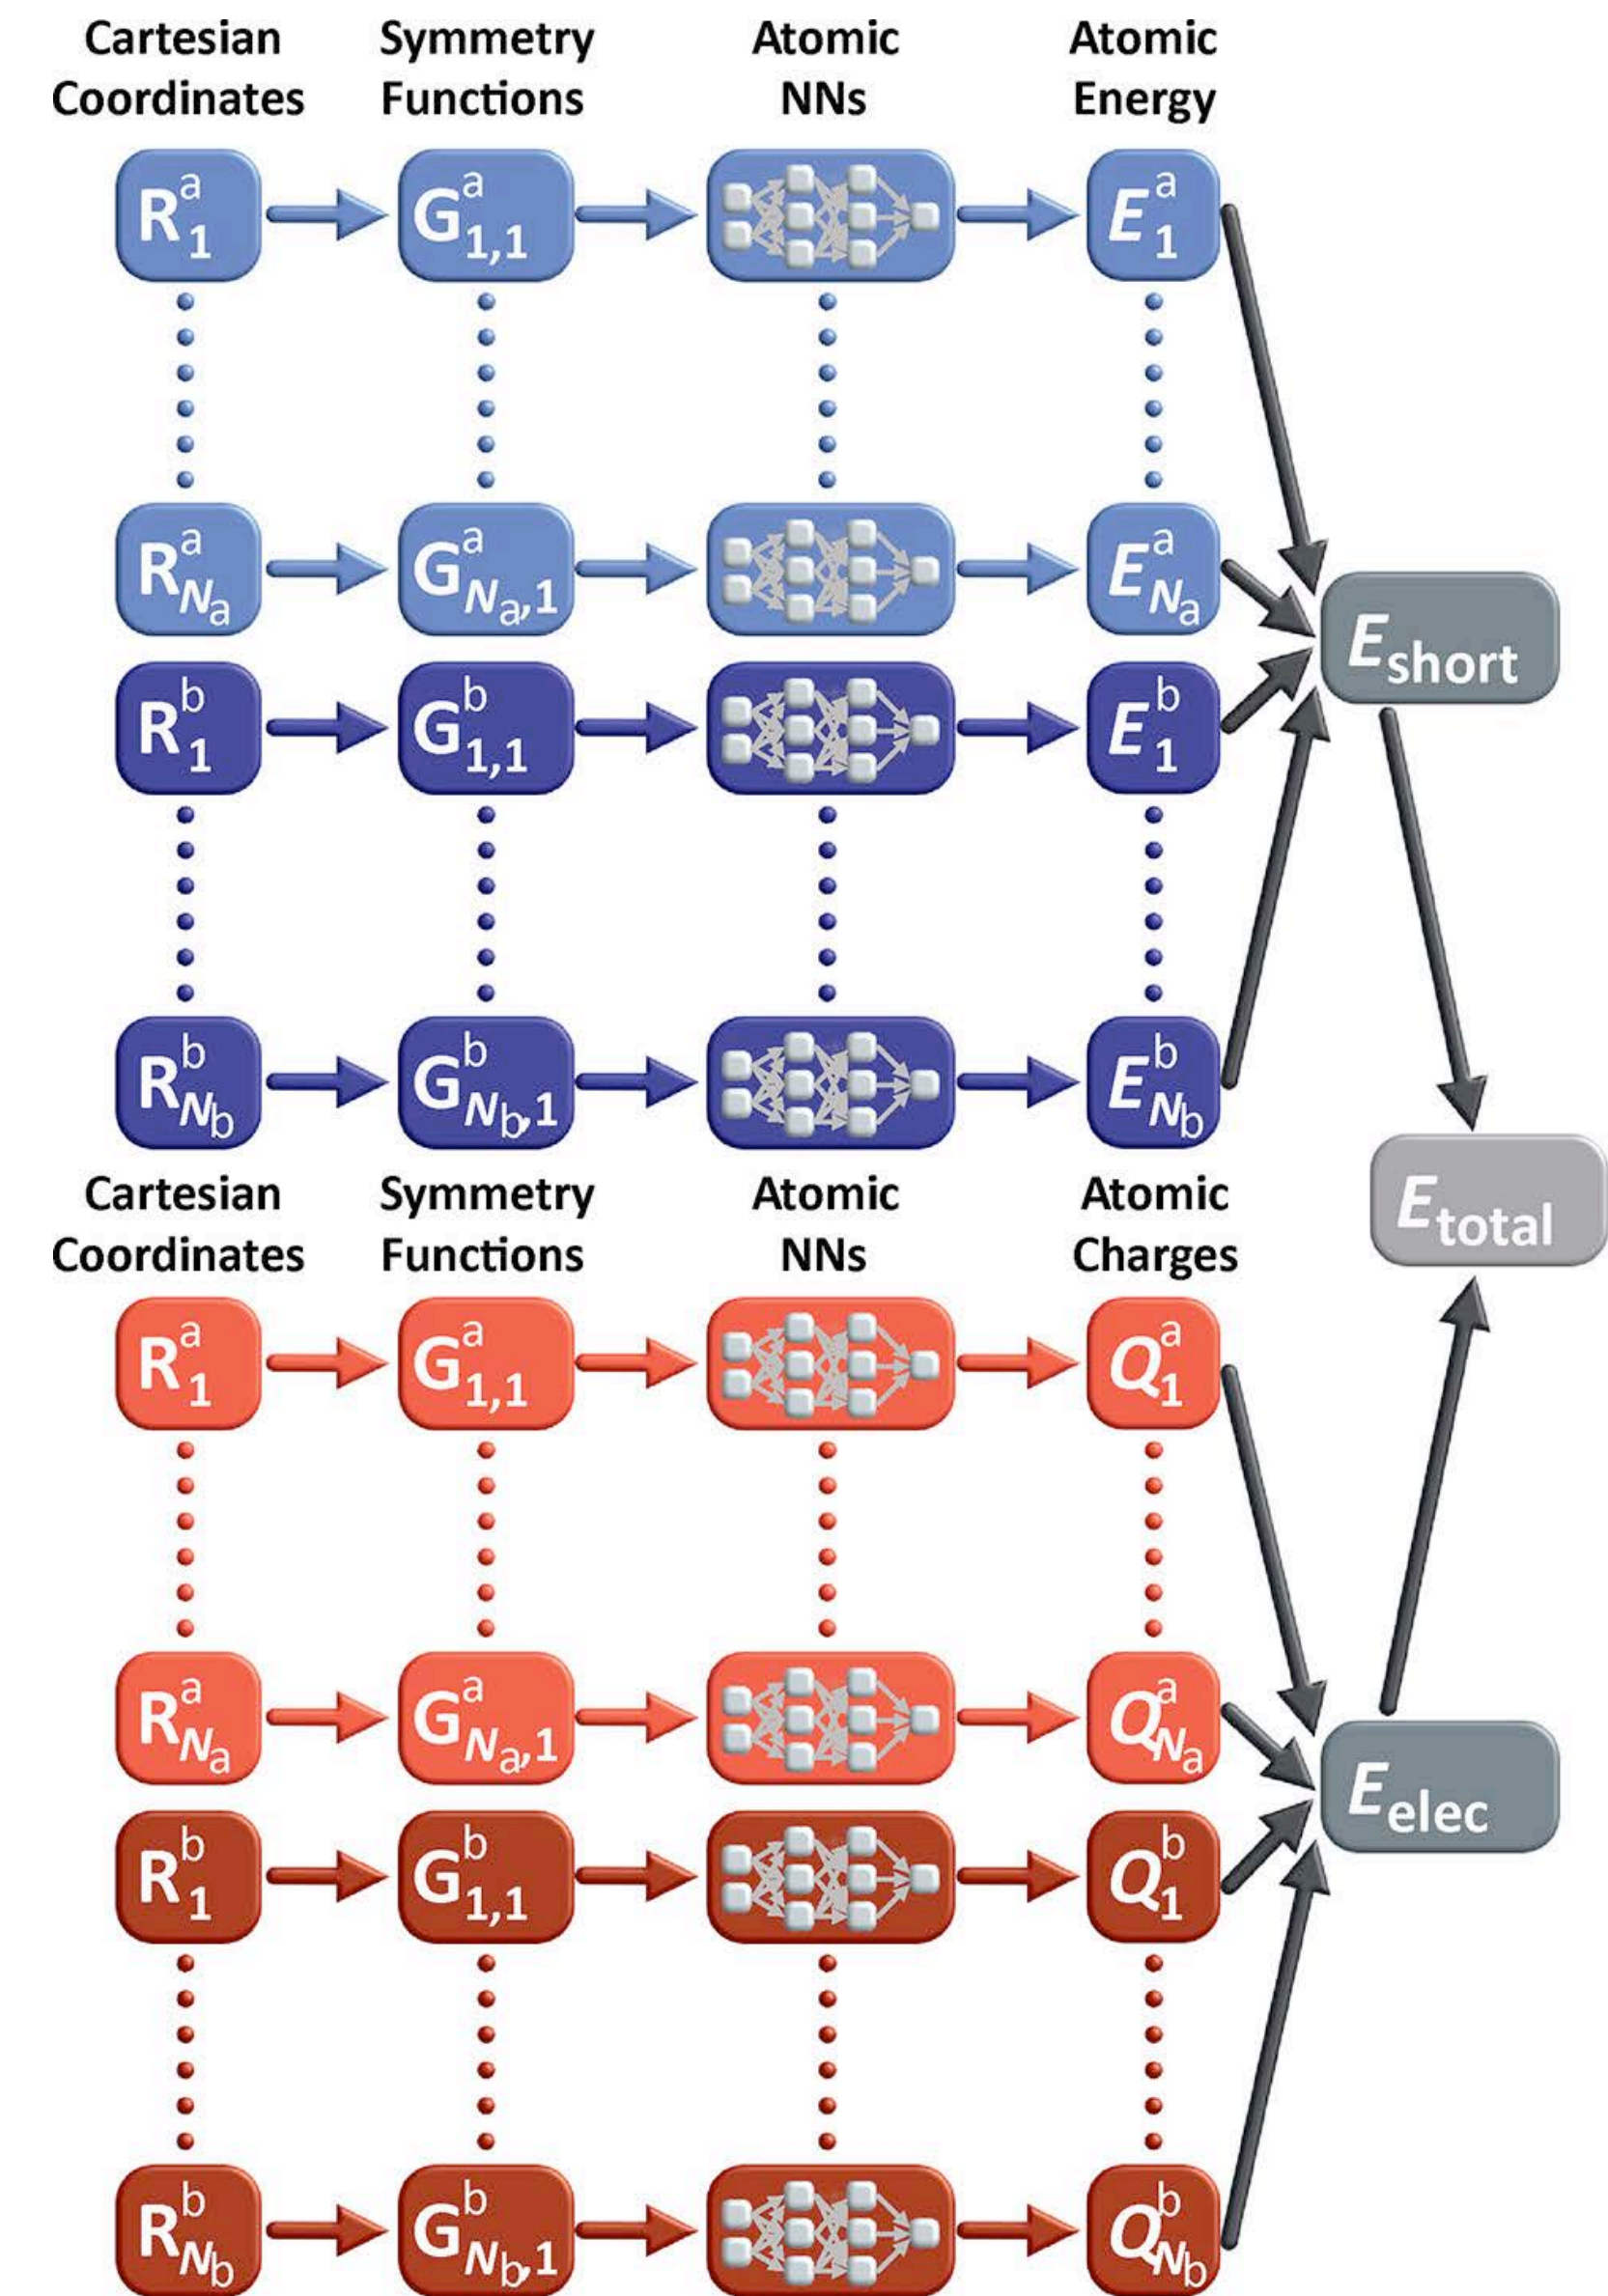
Different approaches have been proposed to deal with long-range interactions

The first option is to include explicit Coulomb terms for fixed charges for each atomic environment:

$$E = \sum_i E_i + E_{\text{elec}}$$

But a better option may be to predict atomic charges using the symmetry functions, then predict a short-range energy and a long-range energy (figure on the right),


$$E = \sum_i E_i + \frac{1}{2} \sum_{i \neq j} \kappa_{ij}(r_{ij}) \frac{Q_i Q_j}{r_{ij}}$$



But there's more: we can use several representations

descriptor	year
atom-centered symmetry functions	2007
bispectrum	2010
Coulomb matrix	2012
SOAP	2013
permutation invariant polynomials	2013
Ewald sum matrix	2015
bag of bonds	2015
overlap matrix	2016
polynomials in MTPs	2016
spherical harmonics	2017
Chebyshev polynomials	2017
many-body tensor representation	2017
histogram of internal coordinates	2017
FCHL	2018
weighted symmetry functions	2018
smoothed atomic densities	2019
orthogonal descriptors	2019
long-distance equivariant repres.	2019

And there is a lot of discussion on why representations matter so much:



(the environments above, for example, are different, but have the same histogram of triangles)

S. Pozdnyakov et al. *PRL* **125**, 166001 (2020)

Furthermore, many-body interactions also matter a lot!

$$E_{tot} = E_{1b} + E_{2b} + E_{3b} + E_{4b} + \dots$$

But there's more: we can use several representations

descriptor	year
atom-centered symmetry functions	2007
bispectrum	2010
Coulomb matrix	2012
SOAP	2013
permutation invariant polynomials	2013
Ewald sum matrix	2015
bag of bonds	2015
overlap matrix	2016
polynomials in MTPs	2016
spherical harmonics	2017
Chebyshev polynomials	2017
many-body tensor representation	2017
histogram of internal coordinates	2017
FCHL	2018
weighted symmetry functions	2018
smoothed atomic densities	2019
orthogonal descriptors	2019
long-distance equivariant repres.	2019

And there is a lot of discussion on why representations matter so much:

What if we are missing critical factors when proposing new descriptors?



The question now becomes: can we do better than designing representations?


Furthermore, many-body interactions also matter a lot!

$$E_{tot} = E_{el} + E_{vib} + E_{rot} + E_{trans} + \dots$$

This is what we will learn next:

Deep Learning

To summarize what we have learned so far

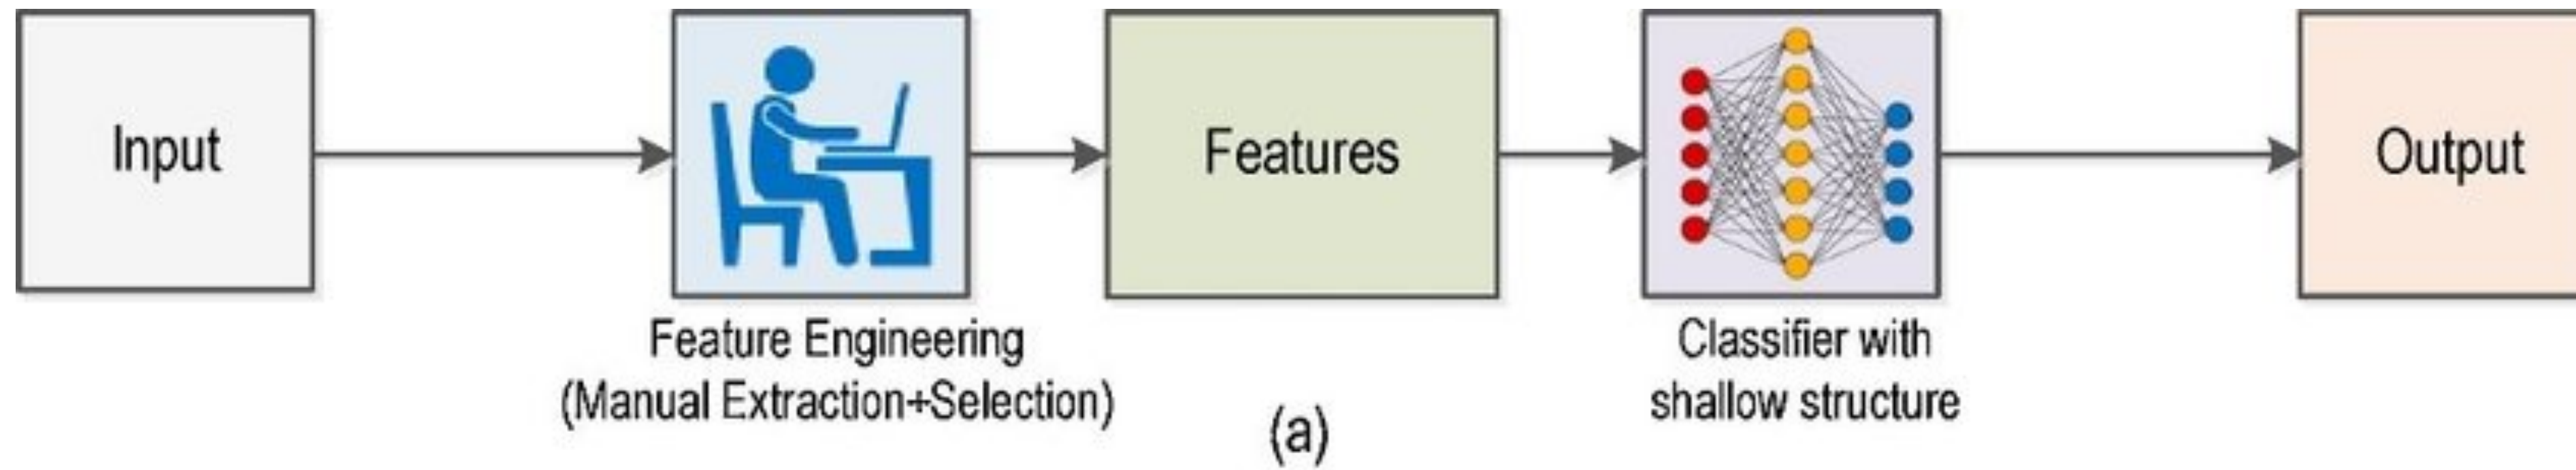


3. Deep Learning Potentials

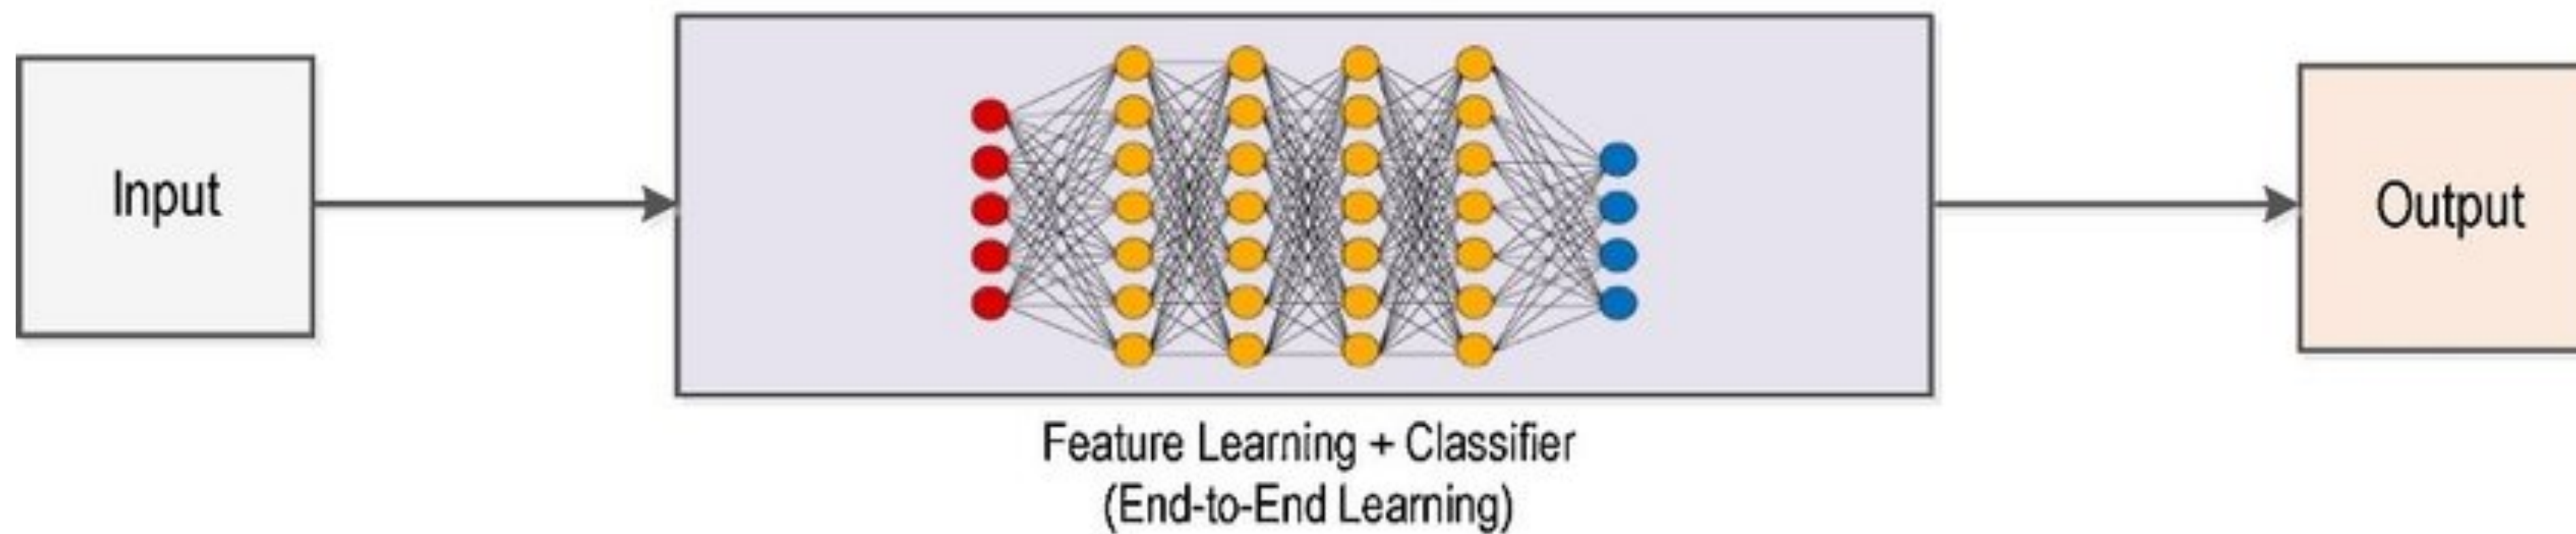
Why use deep learning for interatomic potentials?

A bit of history: why deep learning actually succeeded in other fields?

Traditional ML pipeline



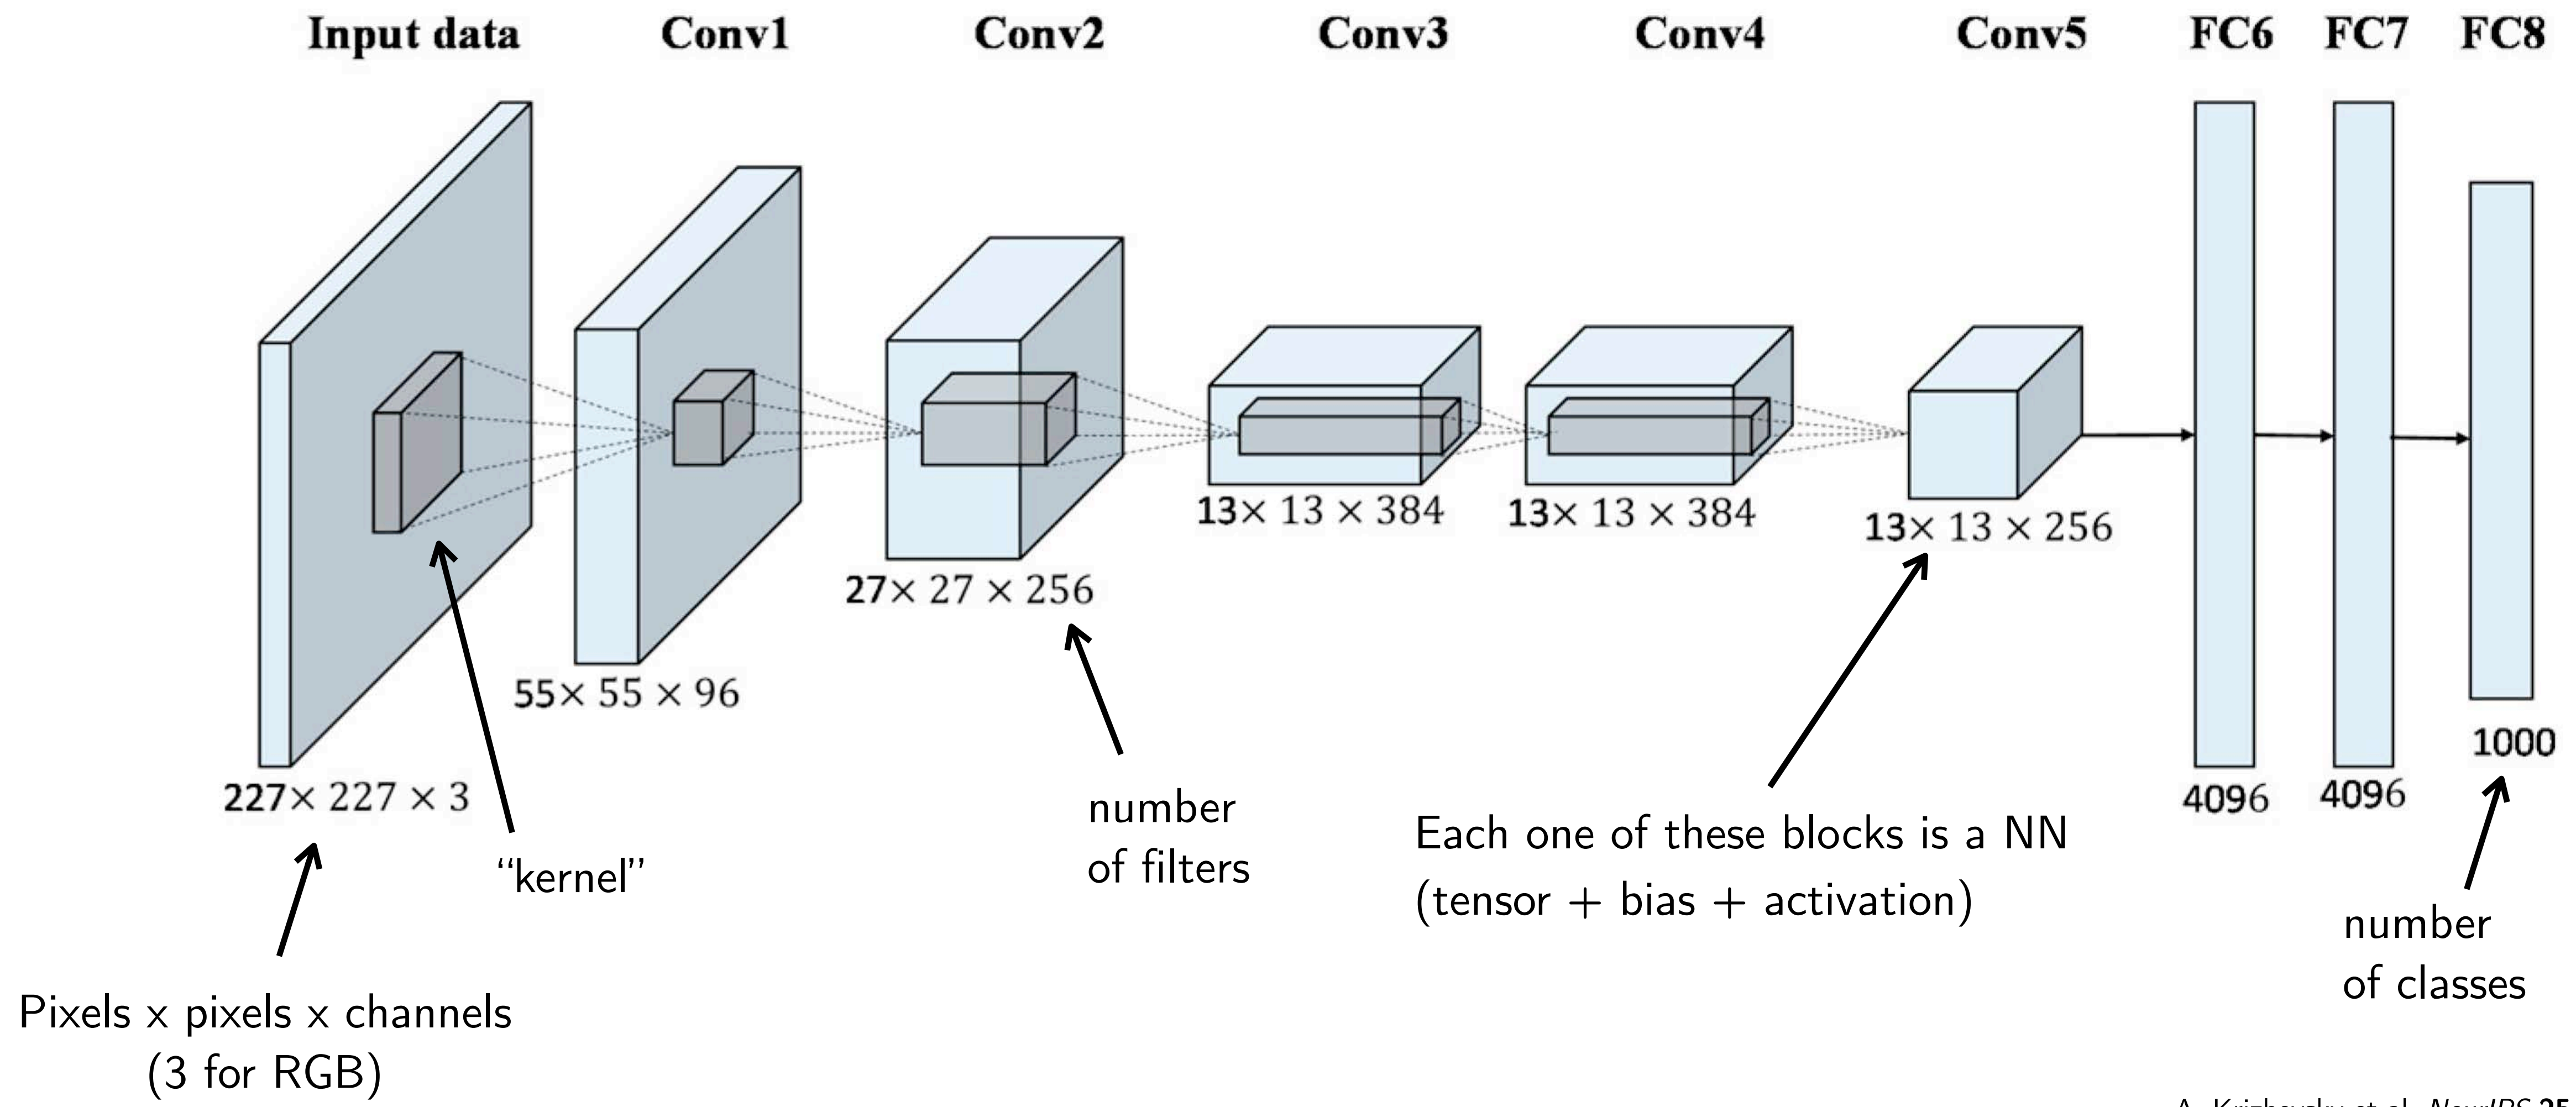
Deep learning pipeline



Deep Learning:
features identified
along with the
training process.

In deep learning, the neural network architectures are more complicated (“deep”)

AlexNet architecture



A. Krizhevsky et al. *NeurIPS* 25 (2012)


Image from X. Han et al. *Remote Sens.* 9 (8), 848 (2017)

If we take each of the filters, we will see the features extracted from the dataset:

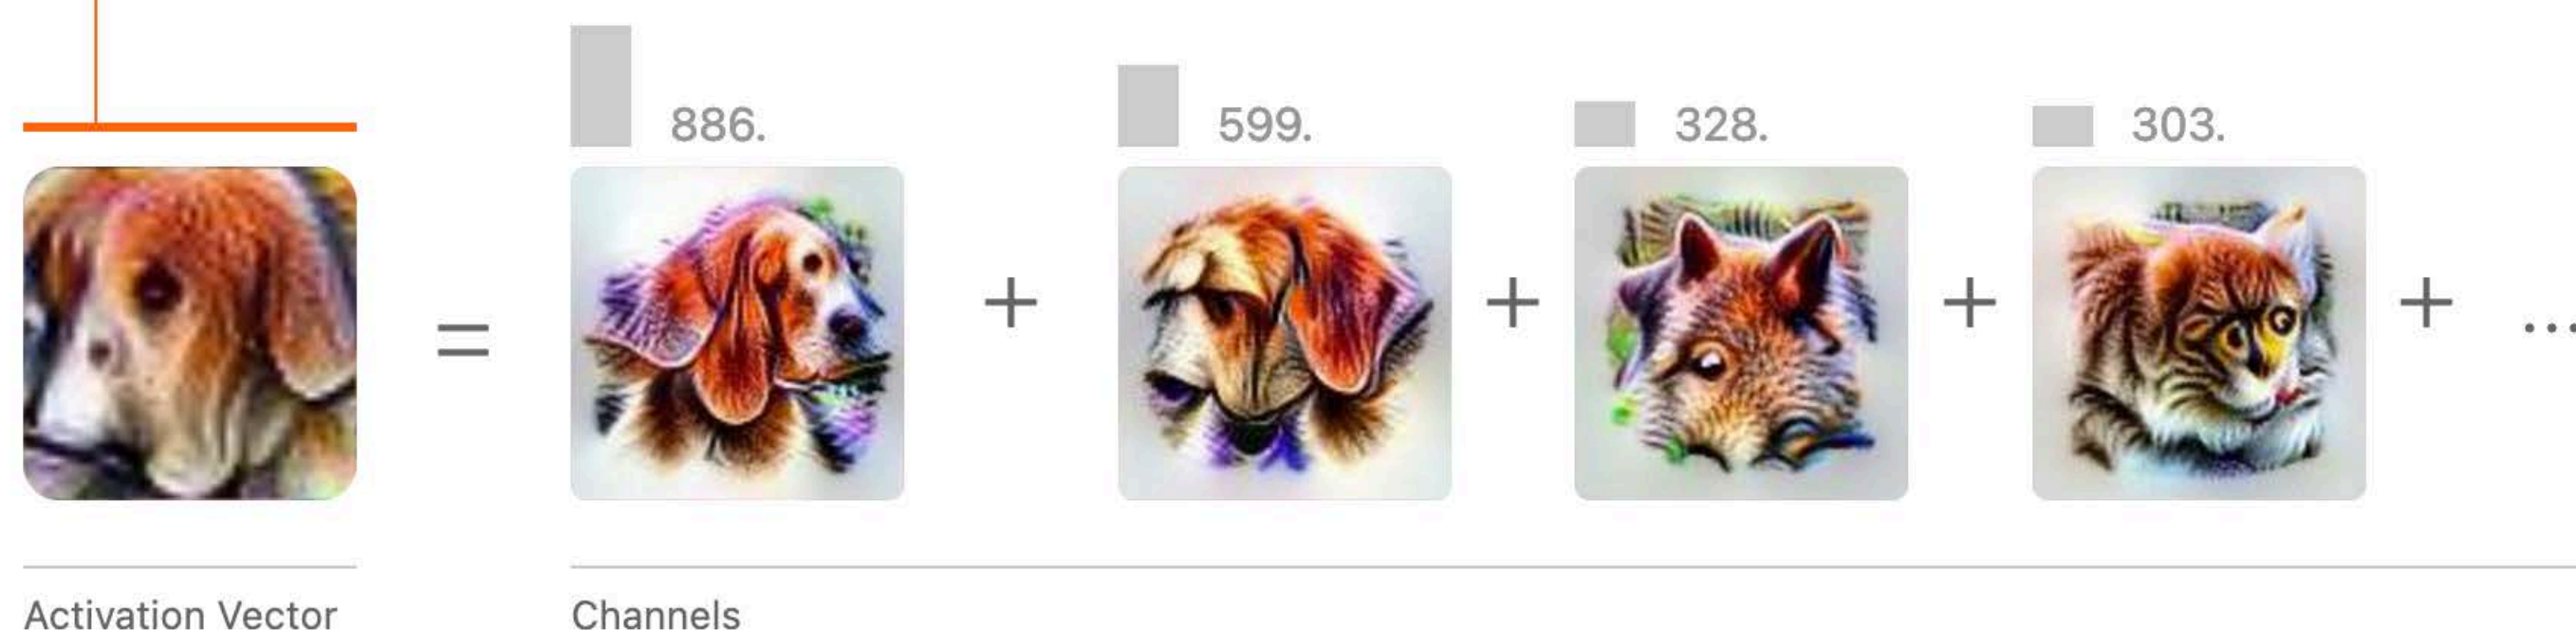


Figure 3: 96 convolutional kernels of size $11 \times 11 \times 3$ learned by the first convolutional layer on the $224 \times 224 \times 3$ input images. The top 48 kernels were learned on GPU 1 while the bottom 48 kernels were learned on GPU 2. See Section 6.1 for details.

What Does the Network See?

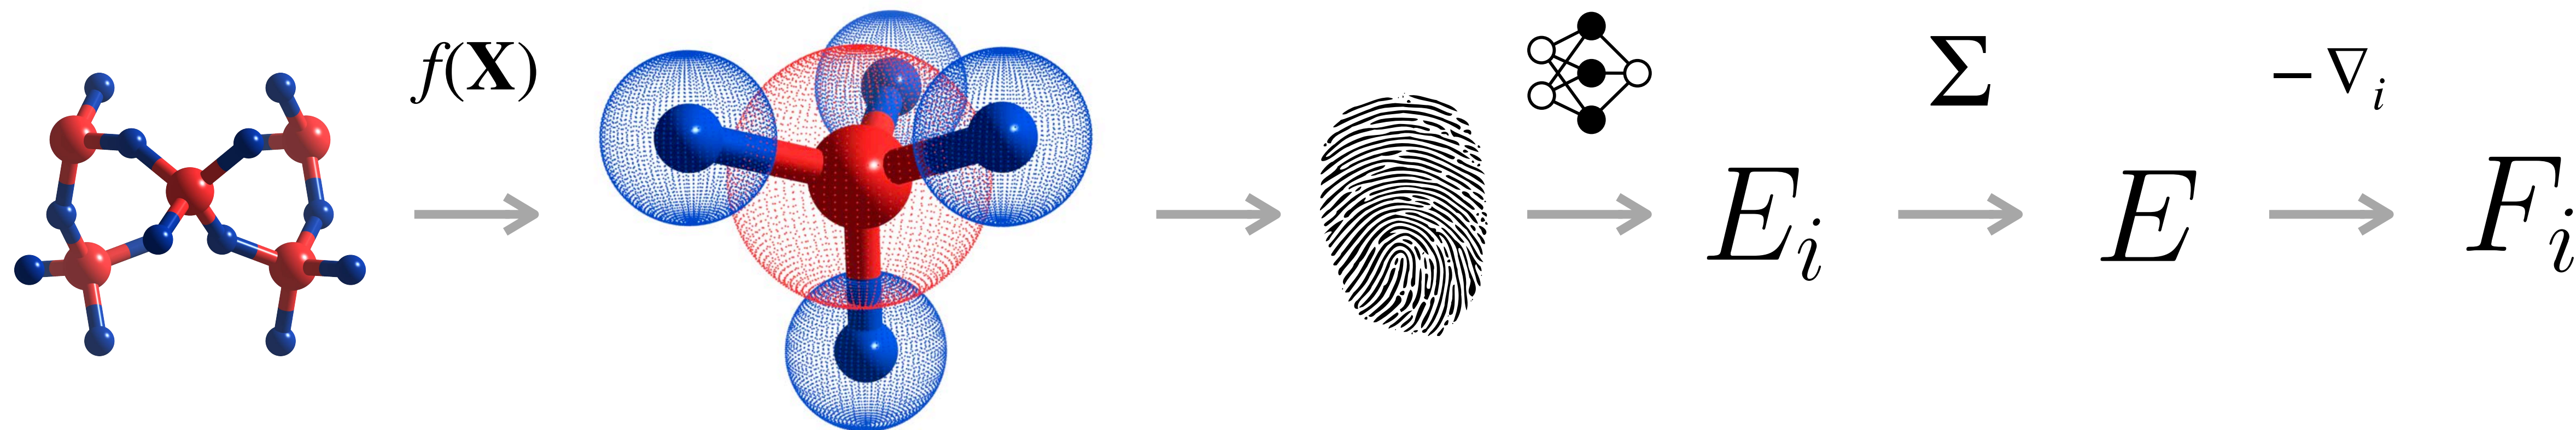


Semantic dictionaries give us a fine-grained look at an activation: what does each single neuron detect? Building off this representation, we can also consider an activation vector as a whole. Instead of visualizing individual neurons, we can instead visualize the *combination* of neurons that fire at a given spatial location. (Concretely, we optimize the image to maximize the dot product of its activations with the original activation vector.)




Mini Tutorial: CNN activations


Featurization Property



Learned Representation Property



But, again, how to represent our data?



In images, a mapping between a $(N \times N \times 3)$ image and (C) classes is well-defined:

$$f : \mathbb{R}^{N \times N \times 3} \rightarrow \mathbb{R}^C$$

(assuming the prediction of class logits)



In materials or molecules, our data is not an image nor a sequence. Instead, our mapping is for a set of atoms in the 3D space:

$$V : \mathbb{Z} \times \mathbb{R}^{N \times 3} \rightarrow \mathbb{R}$$

with the 3D space not well represented by an image.

Why images are not good enough?

A more natural way to represent the data is a graph



By defining a molecule or material system as a graph,

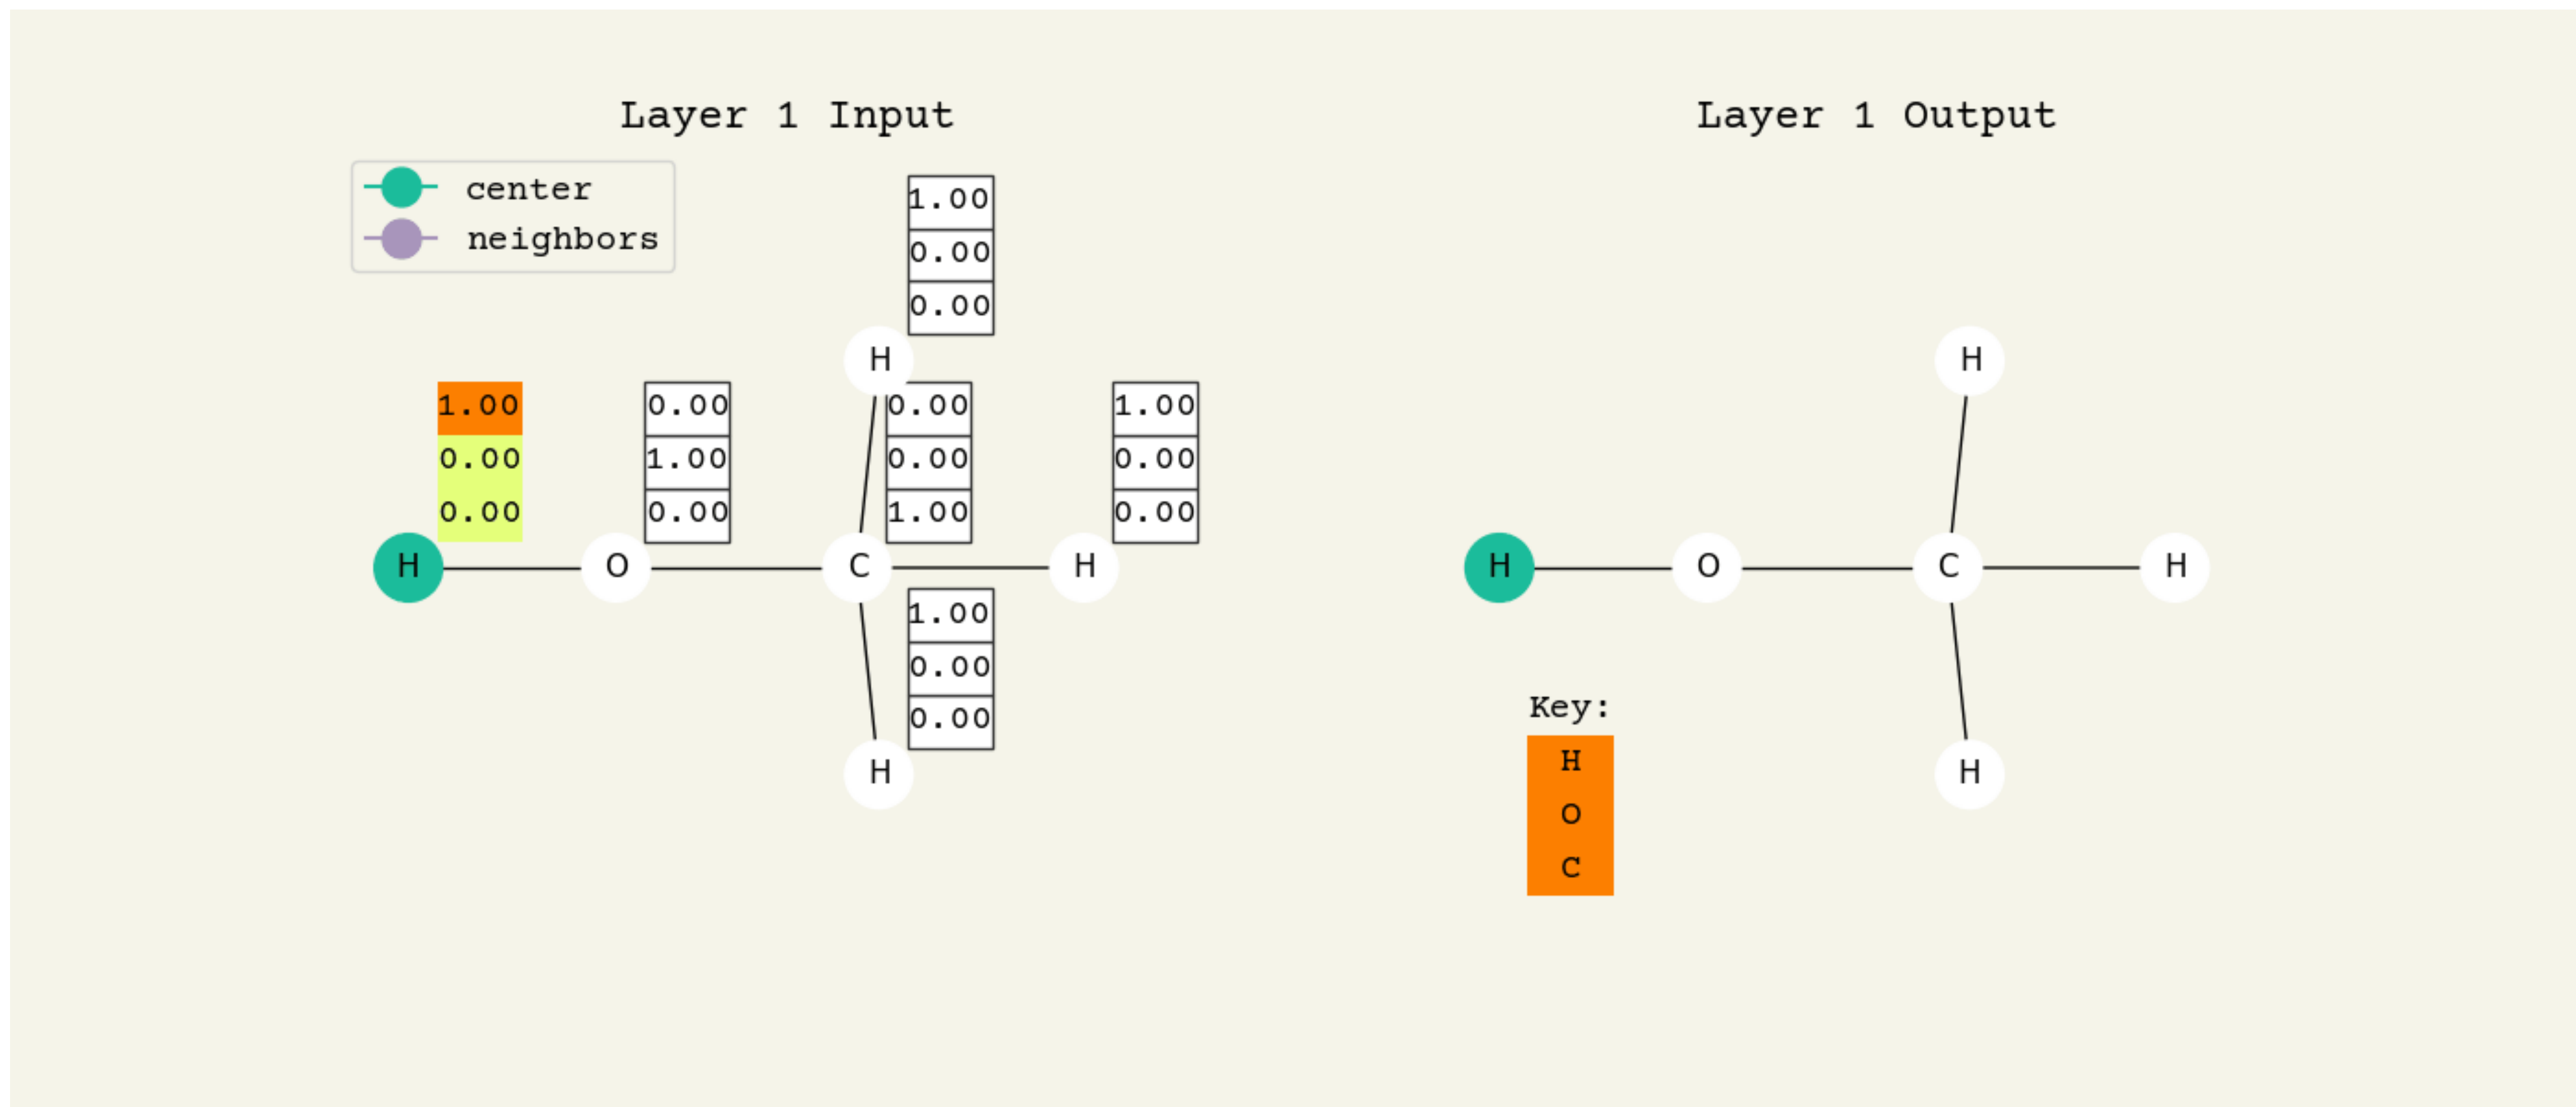
$$G = (V, E)$$

We can initialize some features to each element (or even use random ones):

O	C	H


Now we just need to learn a representation for each atomic environment

How does a graph convolution look like in a molecule?



How does a graph convolution look like in a molecule?

Graph convolution on molecule



In some molecular systems, the graph is simply the connectivity graph of a molecule (covalent bonds).


Each atom interacts with its neighbors, and a “filter” is trained by understanding the interactions between neighbors.

This enables us to predict properties from the molecular graph.



Some problems with a graph-only approach:

1. Connectivity graphs do not tell us anything about conformers (or PESes)
2. What about materials?


Let's examine how the concept of "graph convolution" work with a solid material




Like before, we analyze each atomic environment and create a fingerprint for them



Then, we update the representations using a specific neural network architecture



The whole process is called “message passing” framework




What we are doing is combining the atomic environments in the graph, just like we saw for the CNNs.

The figure on the left shows a central atom and its neighborhood. Atoms which are faded away are not in the neighborhood of this central atom.

A message passing neural network takes the initial graph and representations and creates a node-based representation for each environment.

The architecture of an MPNN



Mathematically, at the layer n for the node i , the message vector $\mathbf{m}_i^{(n+1)}$ is given by

$$\mathbf{m}_i^{(n+1)} = \sum_{j \in \mathcal{N}(i)} M_n \left(\mathbf{h}_i^{(n)}, \mathbf{h}_j^{(n)}, \mathbf{e}_{ij} \right)$$

where $\mathbf{h}_i^{(n)}$ is the representation of node i at layer n , M_n is a neural network, and $\mathcal{N}(i)$ is the neighborhood of i .

The new representation $\mathbf{h}_i^{(n+1)}$ is given by

$$\mathbf{h}_i^{(n+1)} = U_n \left(\mathbf{h}_i^{(n)}, \mathbf{m}_i^{(n+1)} \right)$$

where U_n is a neural network.

This concept of “graph convolution” is very similar to a CNN

CNN on image


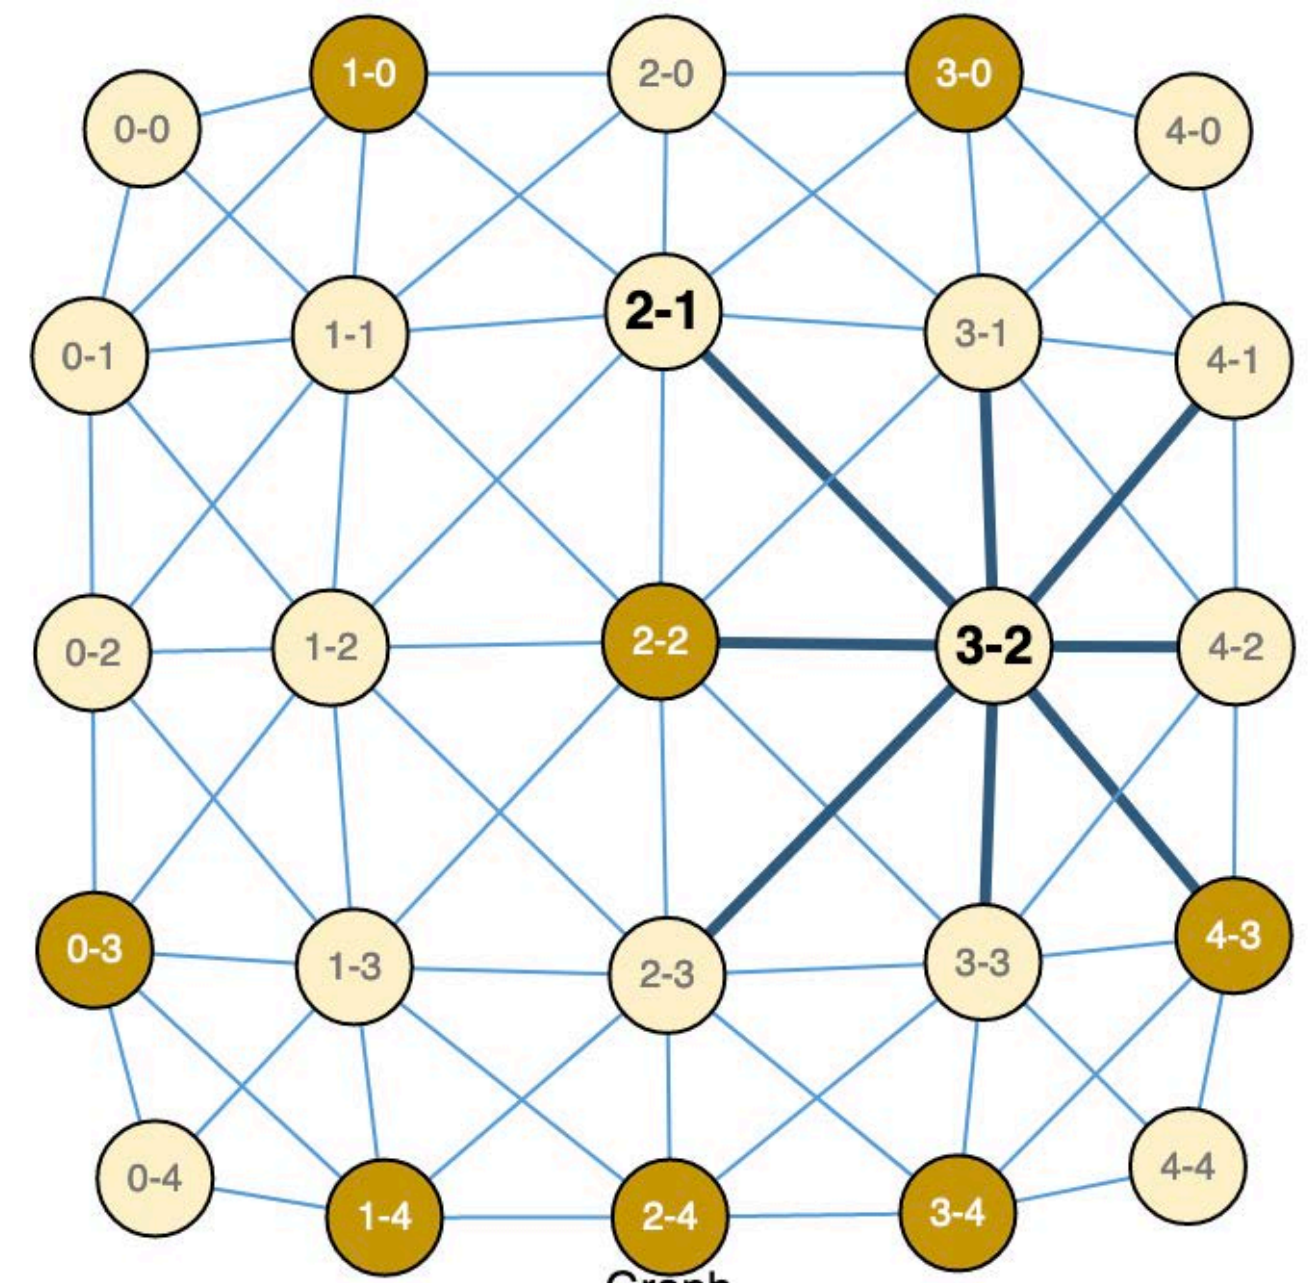



Image Pixels

Graph convolution on image



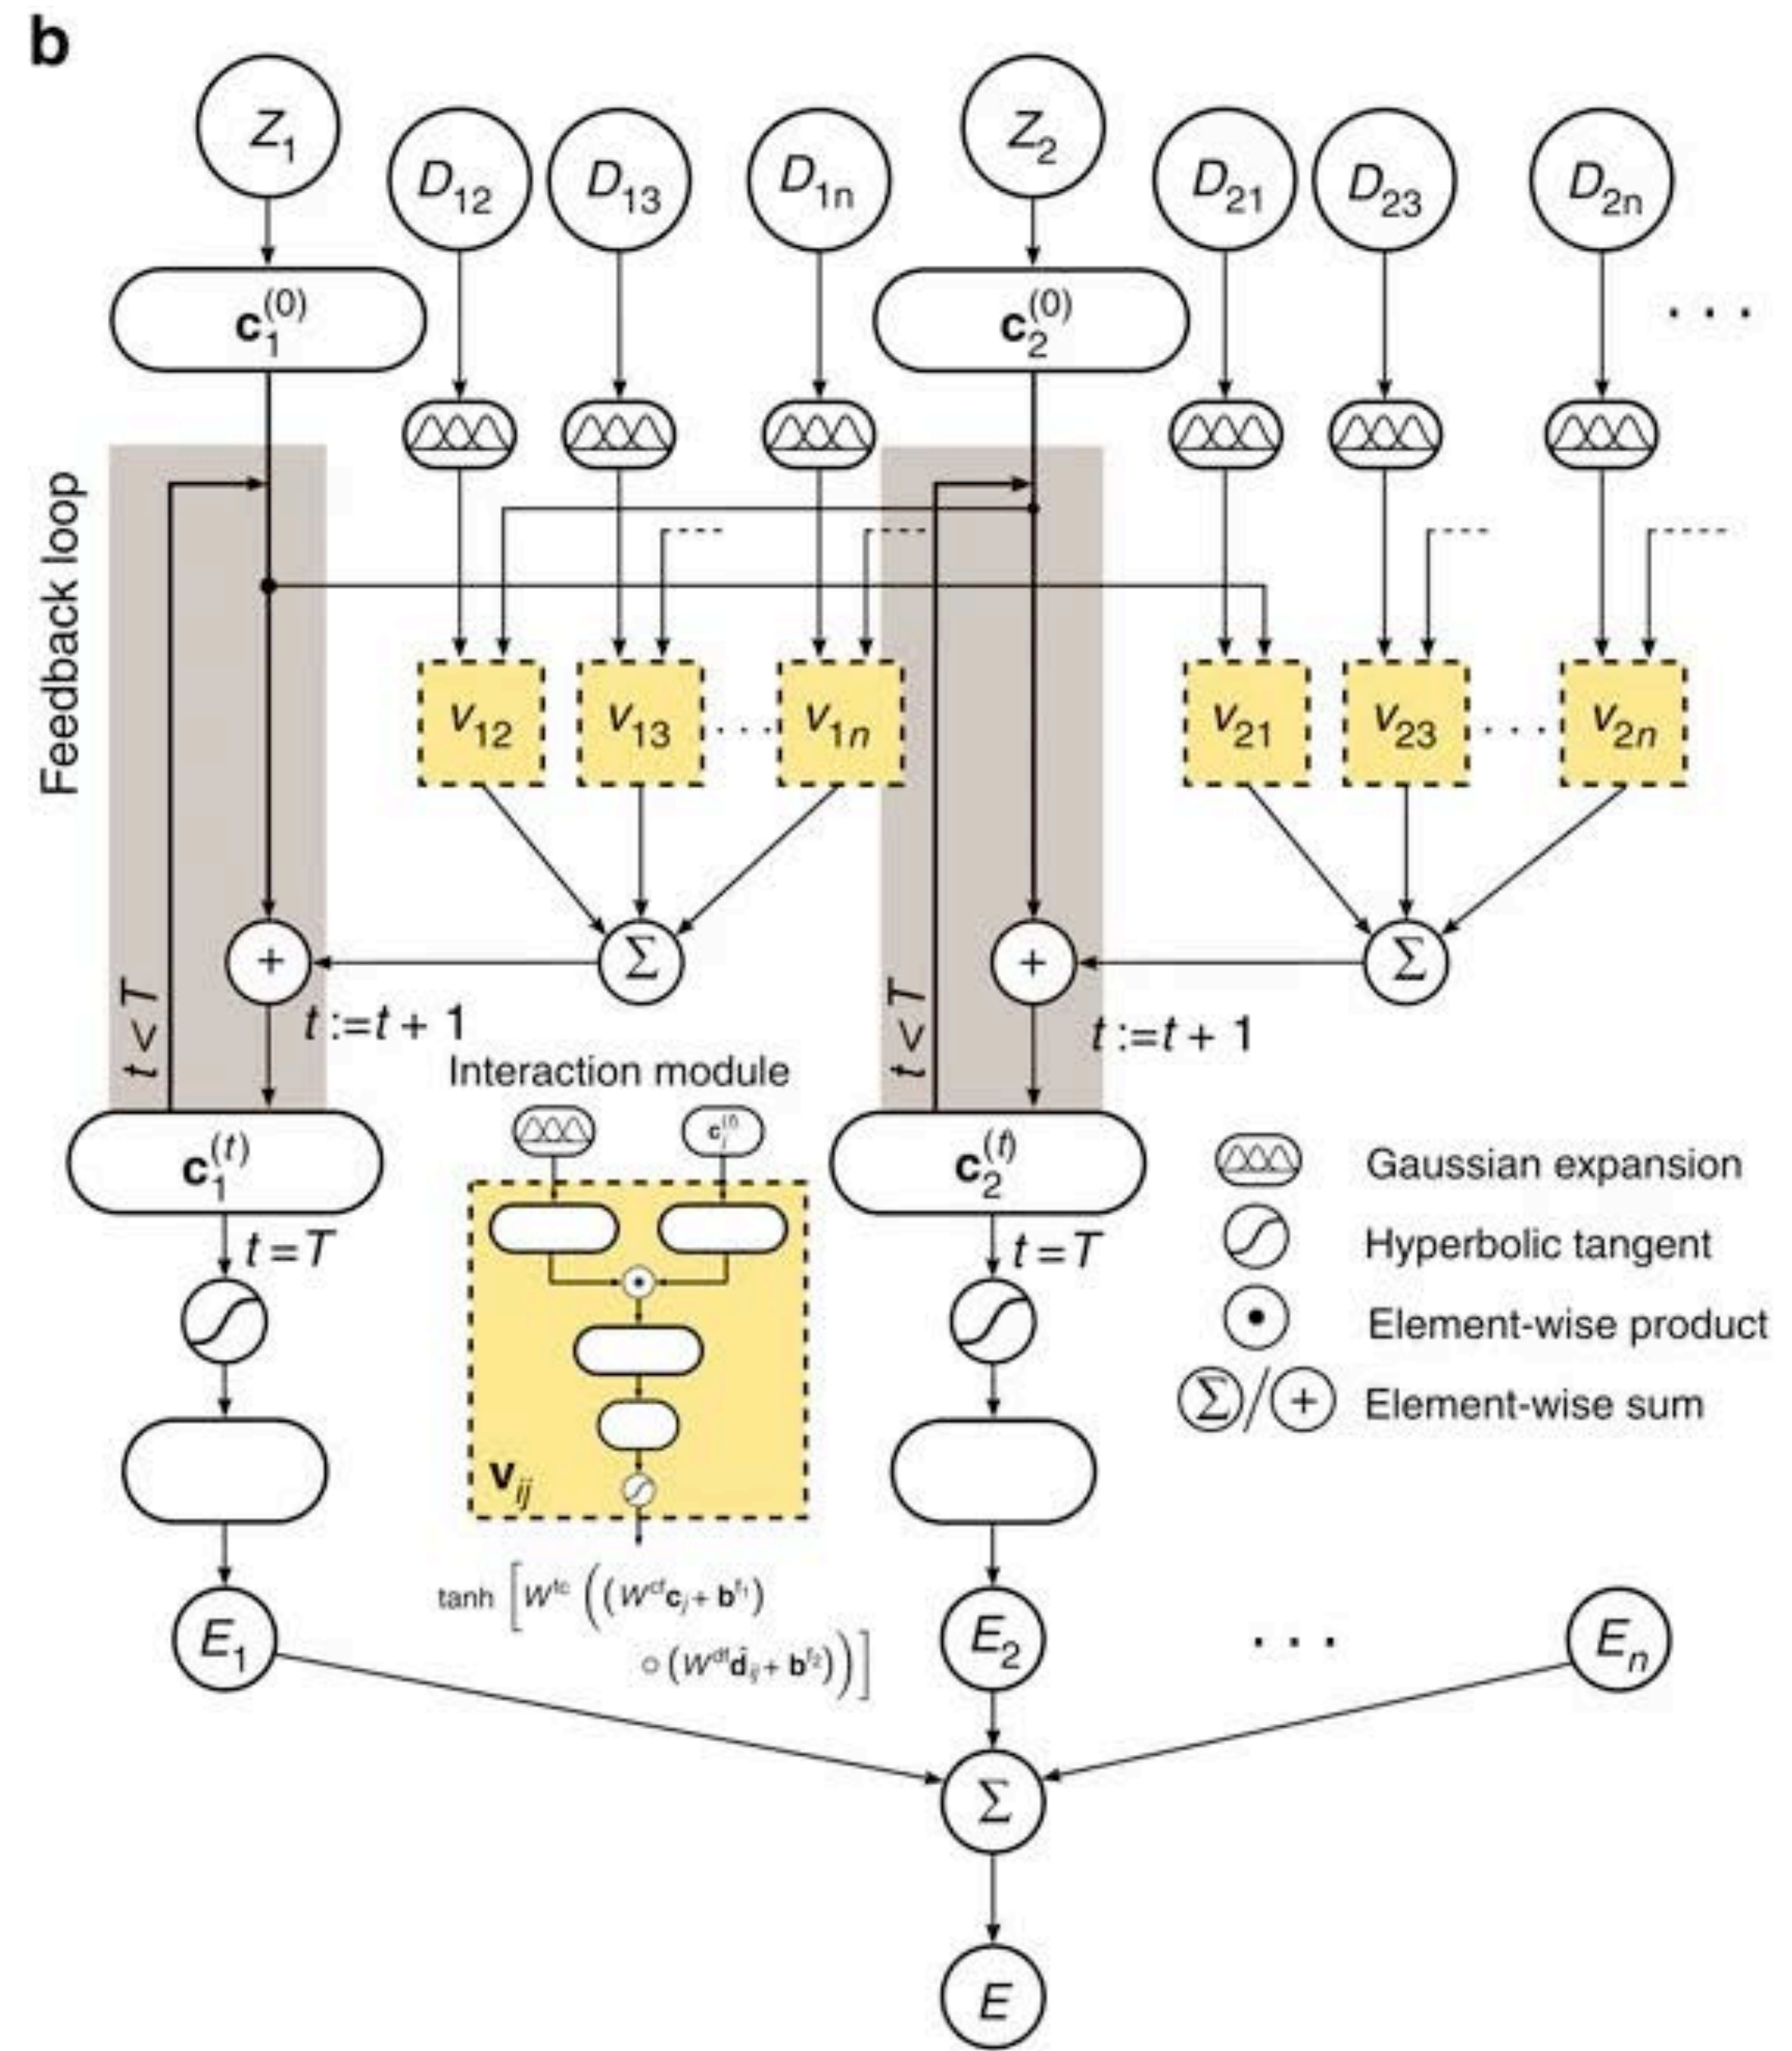
Graph

Graph convolution on geometry

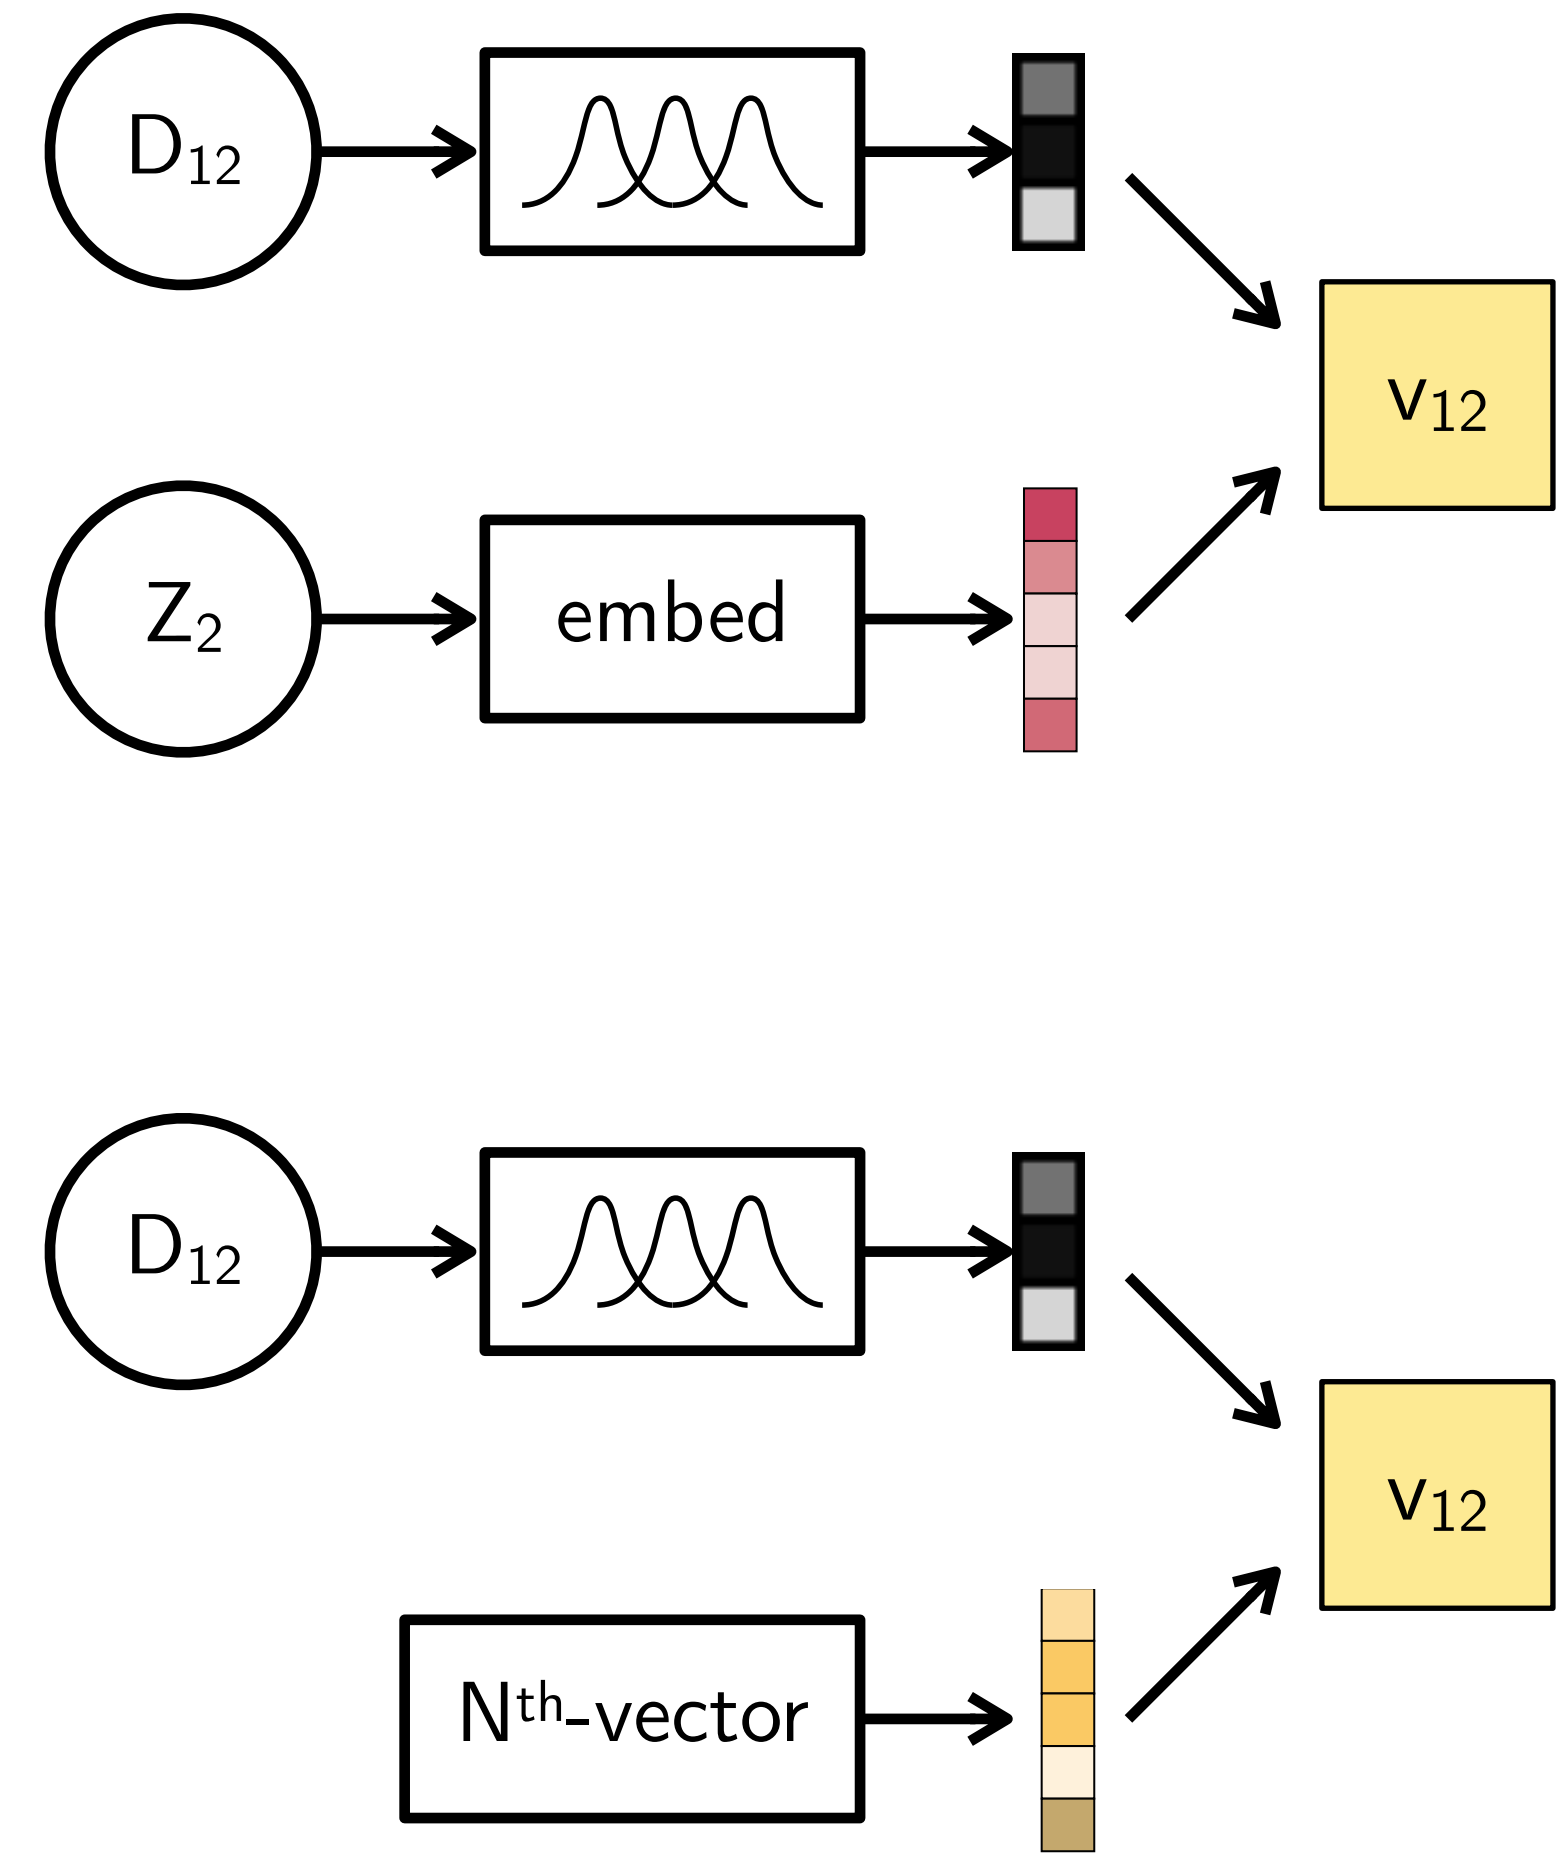


CNN images: B. Sanchez-Lengeling et al. *Distill.pub* (2021)

Application to NNIP: the Deep Tensor Neural Network (DTNN)




The architecture looks complicated! Let's break it down:




DTNN showed interpretable filters and excellent prediction of molecular properties

Local contribution of a test charge
(probing the NN):



Different isomers of $\text{C}_7\text{O}_2\text{H}_{10}$

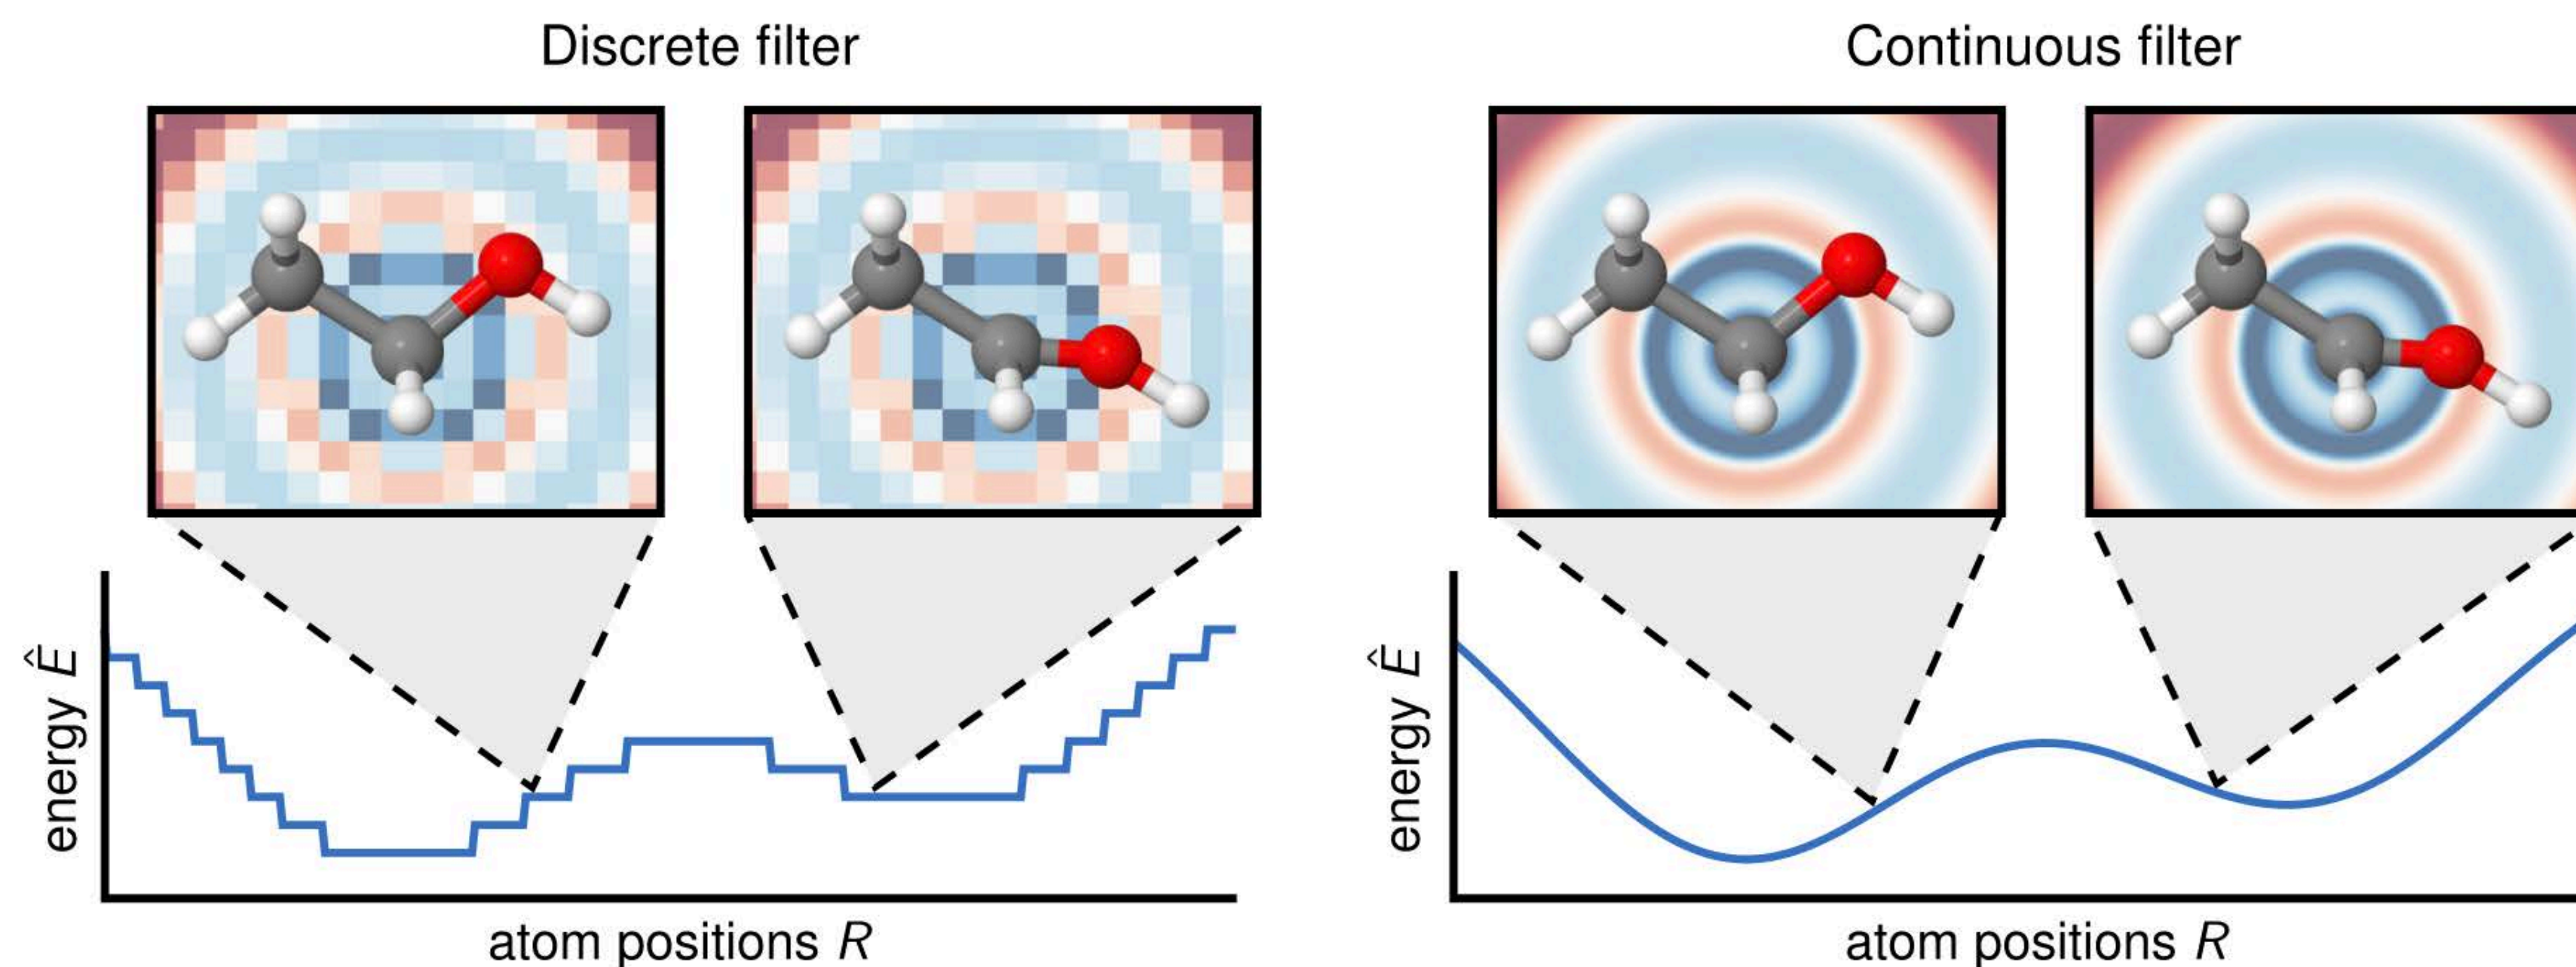


The limitations of the previous models for performing simulations at $T > 0$ K

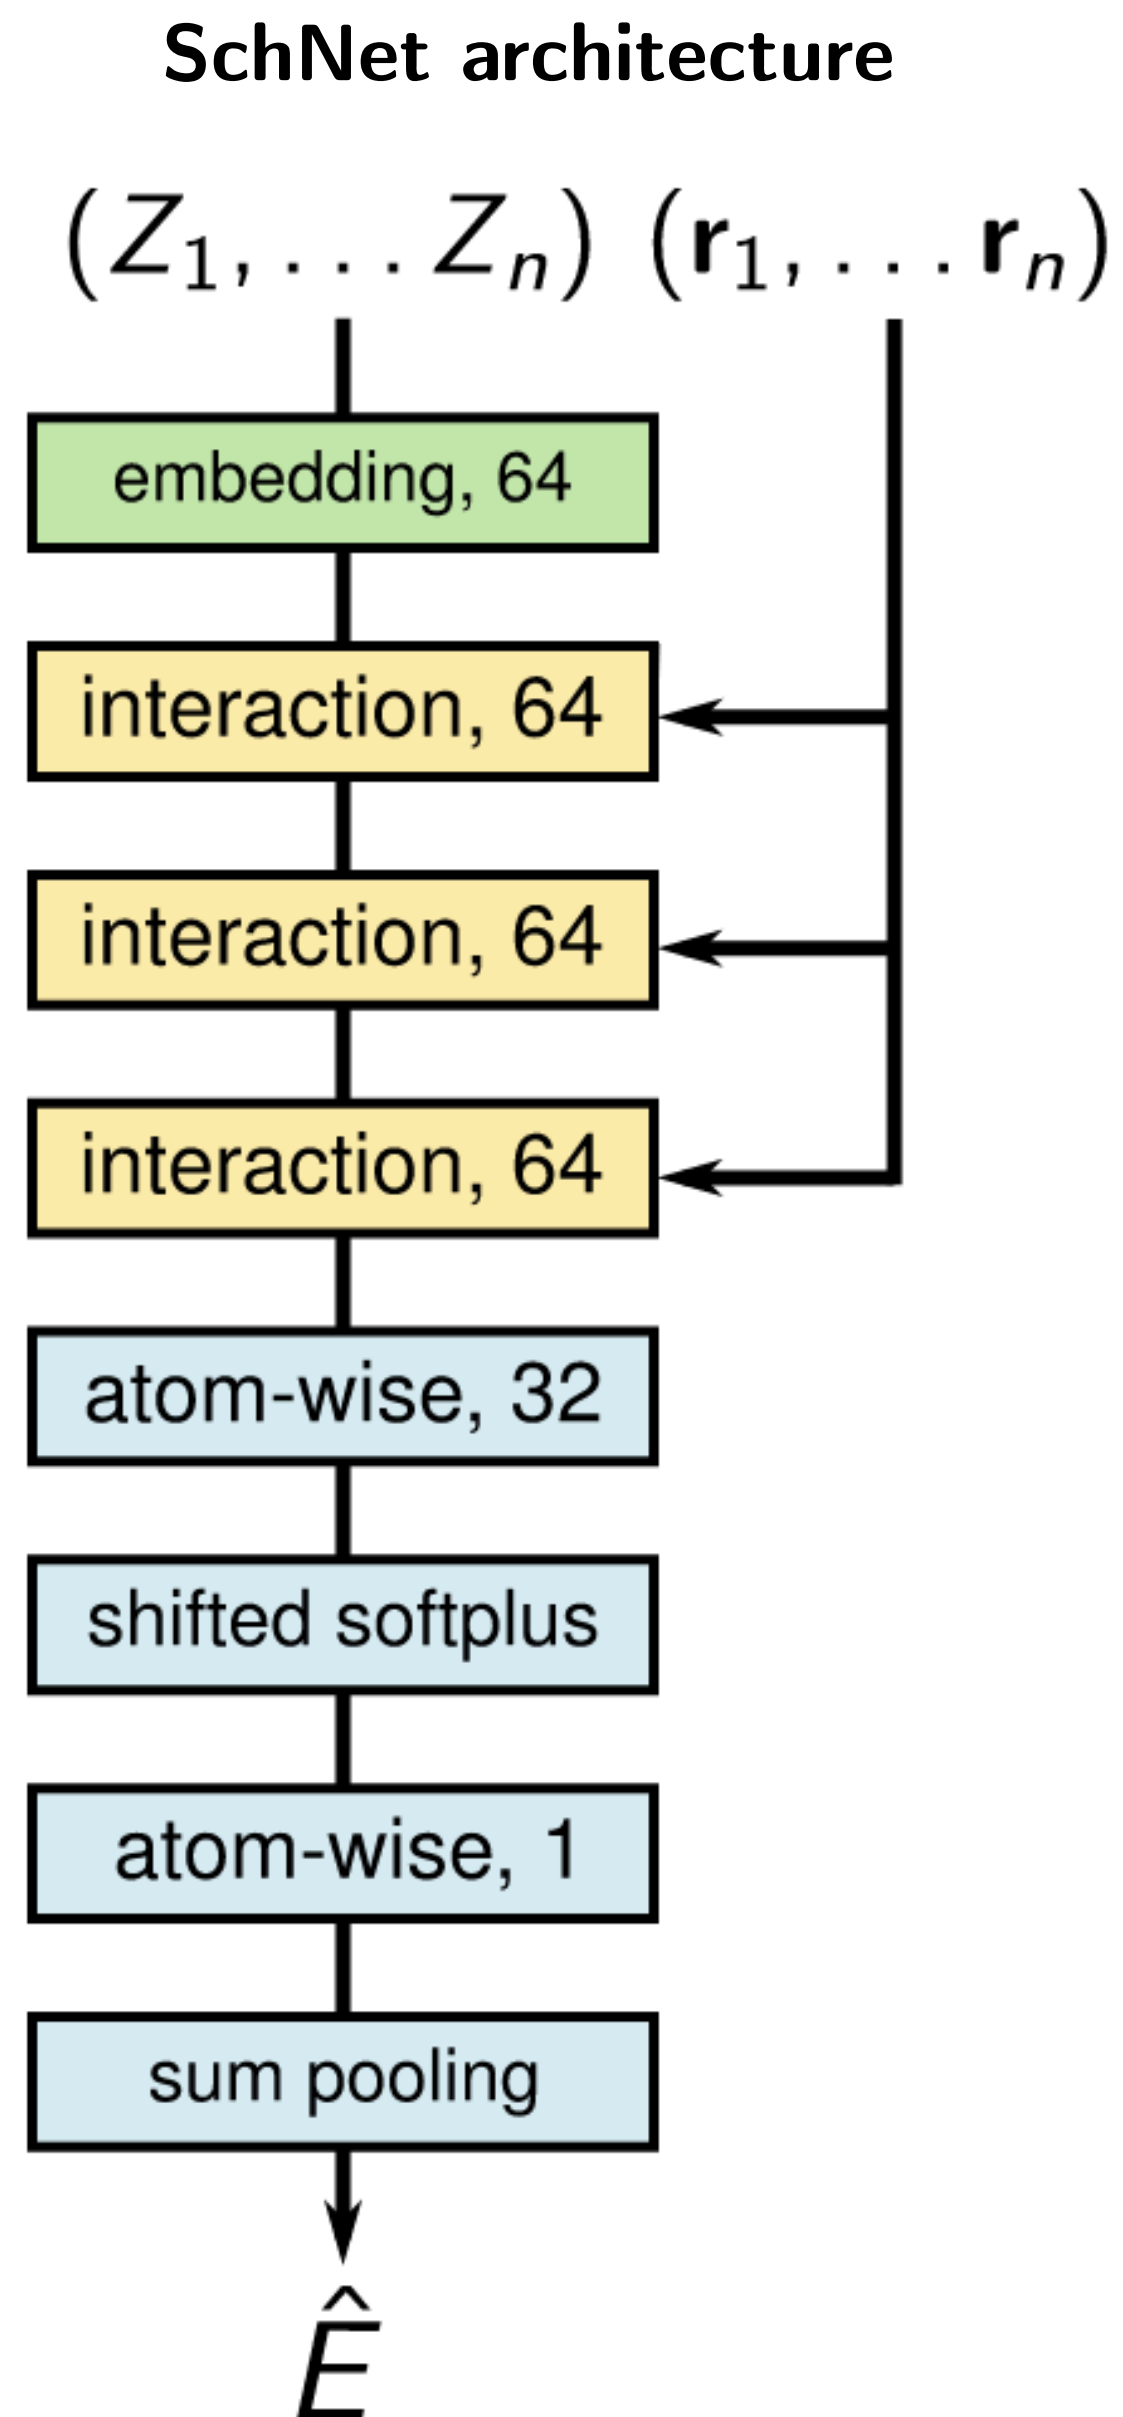
So far, the graph-based NNs have treated molecular graphs or ground-state properties. What if we wanted to perform MD simulations?

$$\longrightarrow \mathbf{F} = -\nabla_{\mathbf{r}} E$$

The problem with the previous NN architectures is that they mostly predict properties of a static graph or 3D structure. If the atoms move, it is not guaranteed to vary the energy continuously.



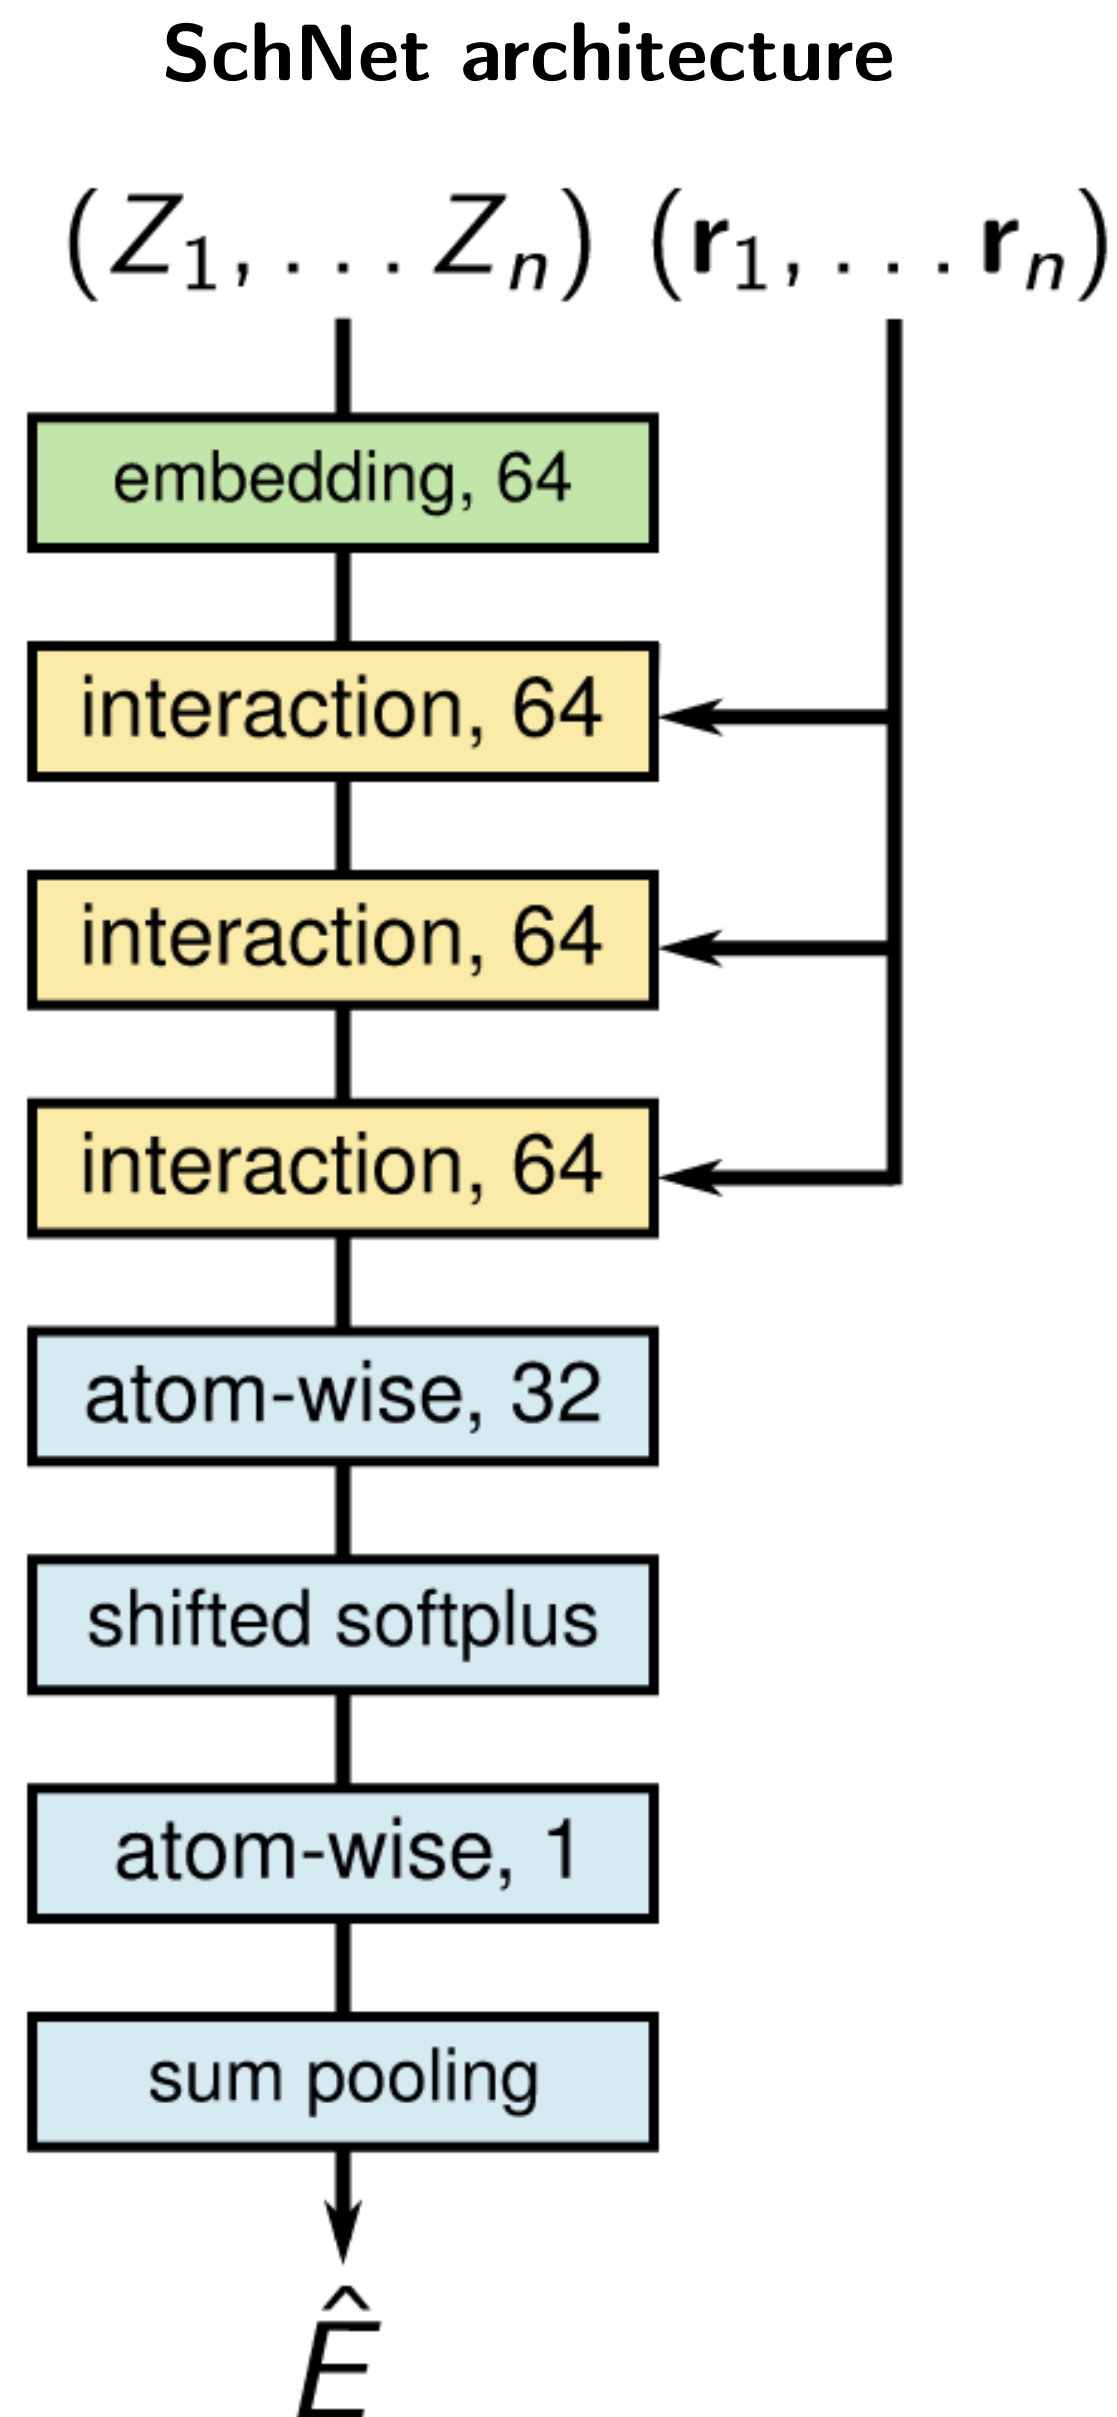
SchNet as a continuous-filter NNIP



The SchNet architecture is not too different from what we learned. On the image on the left, we can see:

1. **embedding layers**, that map an atomic number to a vector.
2. **interaction blocks** for representation learning with message passing.
3. **fully connected NNs**, for predicting atom-wise energies (or properties).
4. **a sum at the end**, pooling all the contributions from the atoms of the systems.

How to fit to energies and forces?



The trick to training this NN is to use information not only about the energy of the PES, but also the forces:

$$\mathcal{L} = \lambda_E \|E - \hat{E}\|^2 + \lambda_F \frac{1}{n} \sum_{i=0}^n \|\mathbf{F}_i - \hat{\mathbf{F}}_i\|^2$$

Where the predicted forces can be obtained by differentiating the predicting energy with respect to the input coordinates:

$$\hat{\mathbf{F}}_i = - \frac{\partial \hat{E}}{\partial \mathbf{R}_i}$$

How does it perform?


Predictions of energies + forces is better than just energies.

Why do you think this is the case?

Predictions of forces enables other properties to be obtained, such as vibrational spectra.


Table 3: Mean absolute errors on $C_7O_2H_{10}$ isomers in kcal/mol.

		mean predictor	SchNet	
			<i>energy</i>	<i>energy+forces</i>
known molecules / unknown conformation	<i>energy</i>	14.89	0.52	0.36
	<i>forces</i>	19.56	4.13	1.00
unknown molecules / unknown conformation	<i>energy</i>	15.54	3.11	2.40
	<i>forces</i>	19.15	5.71	2.18



Many more models were proposed in the field

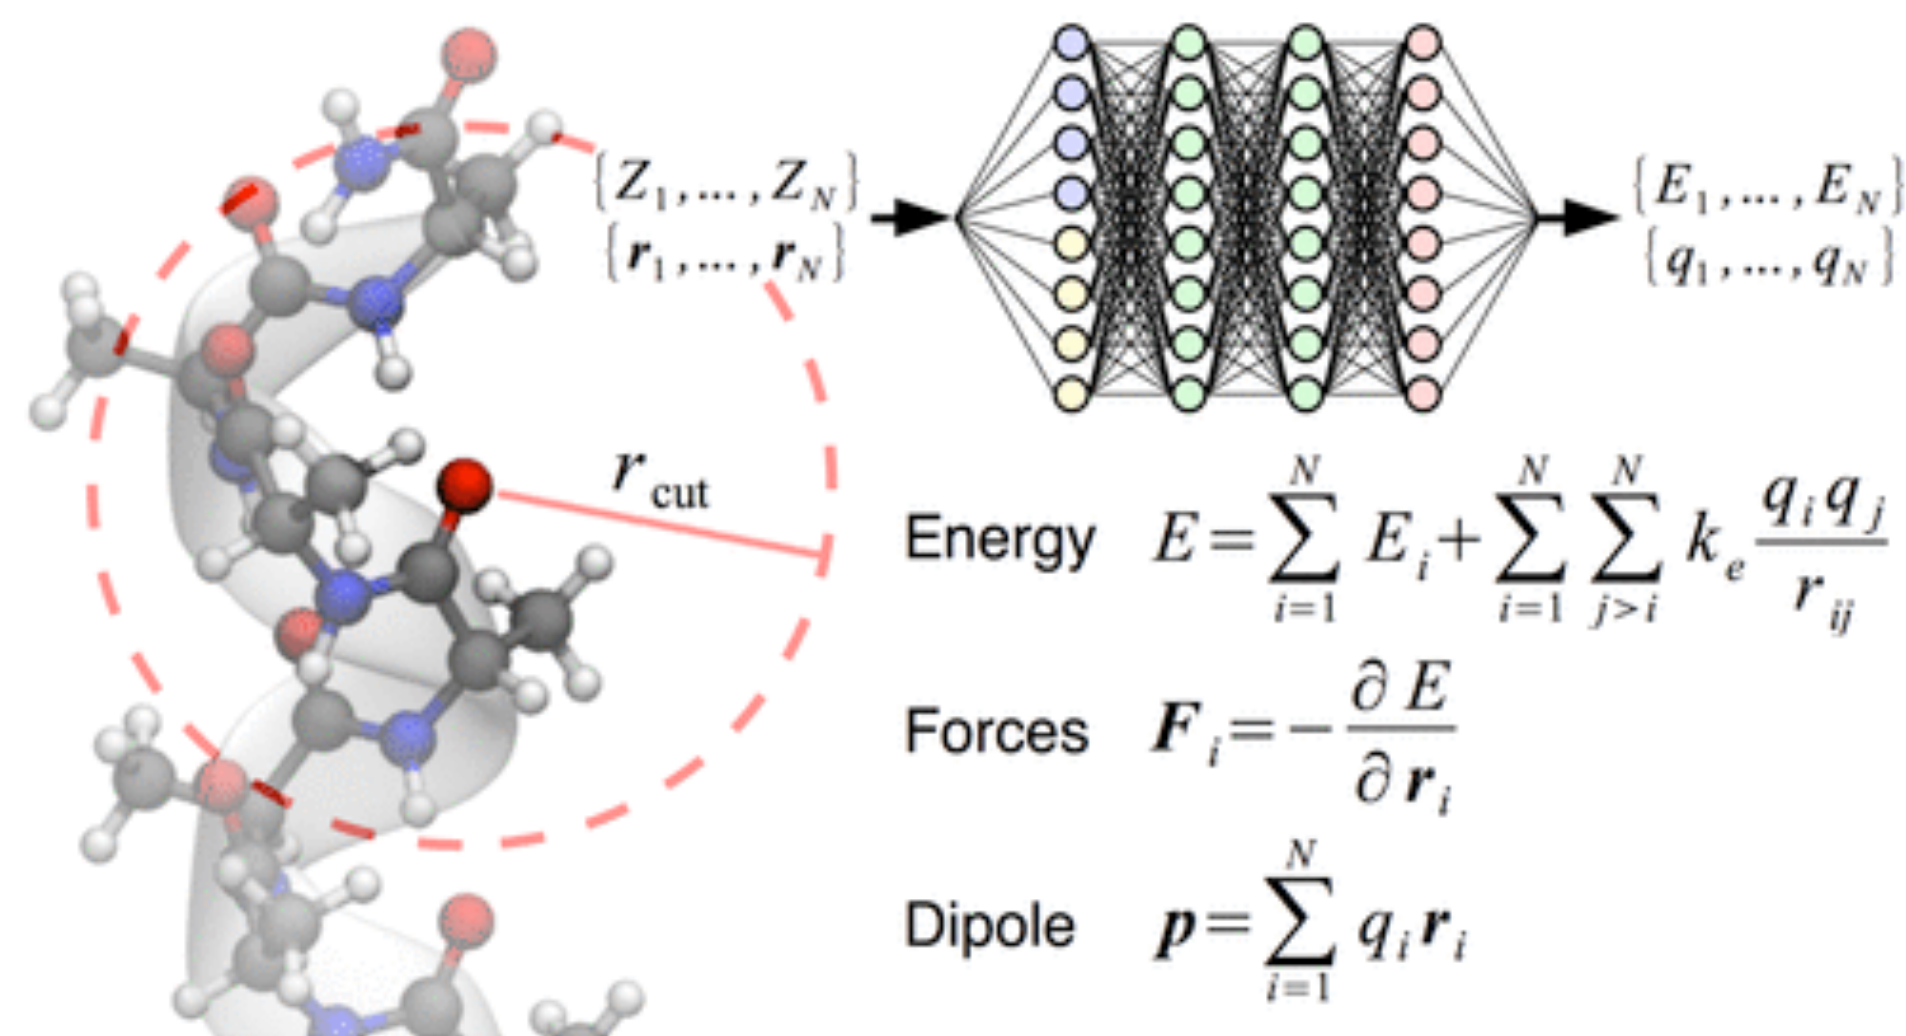
HIP-NN



N. Lubbers et al. *J. Chem. Phys* **148**, 241715 (2018)

Architecture inspired in many-body expansions


PhysNet



O. Unke and M. Meuwly. *JCTC* **15** (6), 3678 (2019)

Prediction of atomic charges

DimeNet



J. Gastegger et al. *ICLR* (2020), *arXiv:2003.03123*

Explicit treatment of three-body terms


There are many more models in the field nowadays. What is being improved?

First improvement: symmetry, invariance, and equivariance

Images are invariant to translation, mirror (often), and rotation (sometimes).

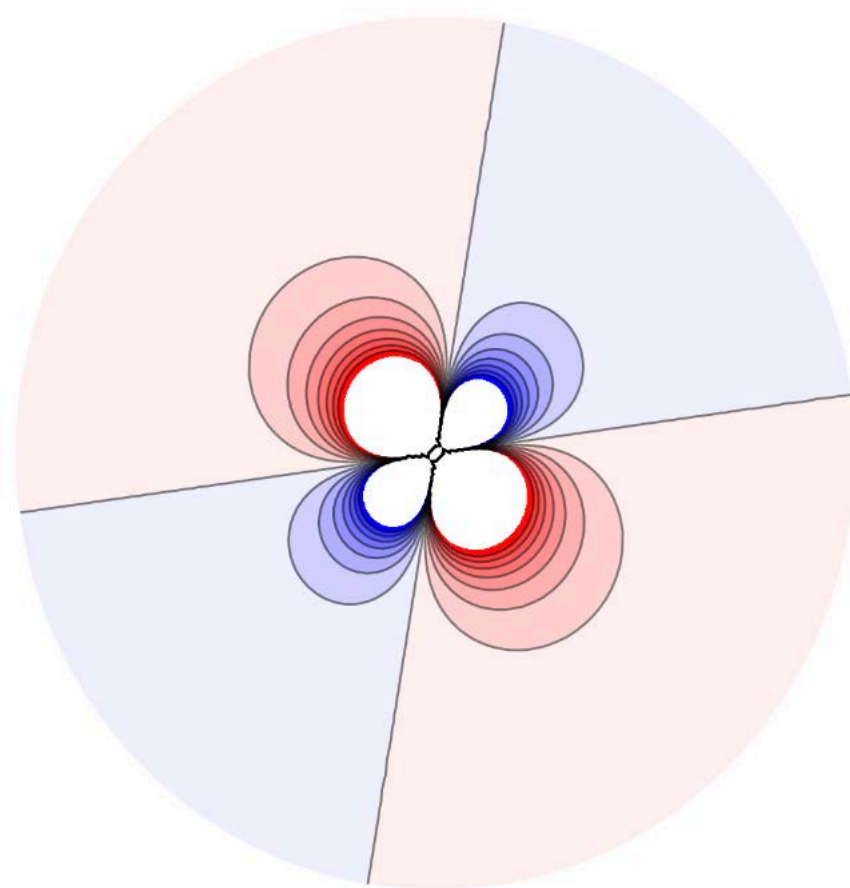


Forces are equivariant to rotation: they transform according to the operation



Recent improvements in this area

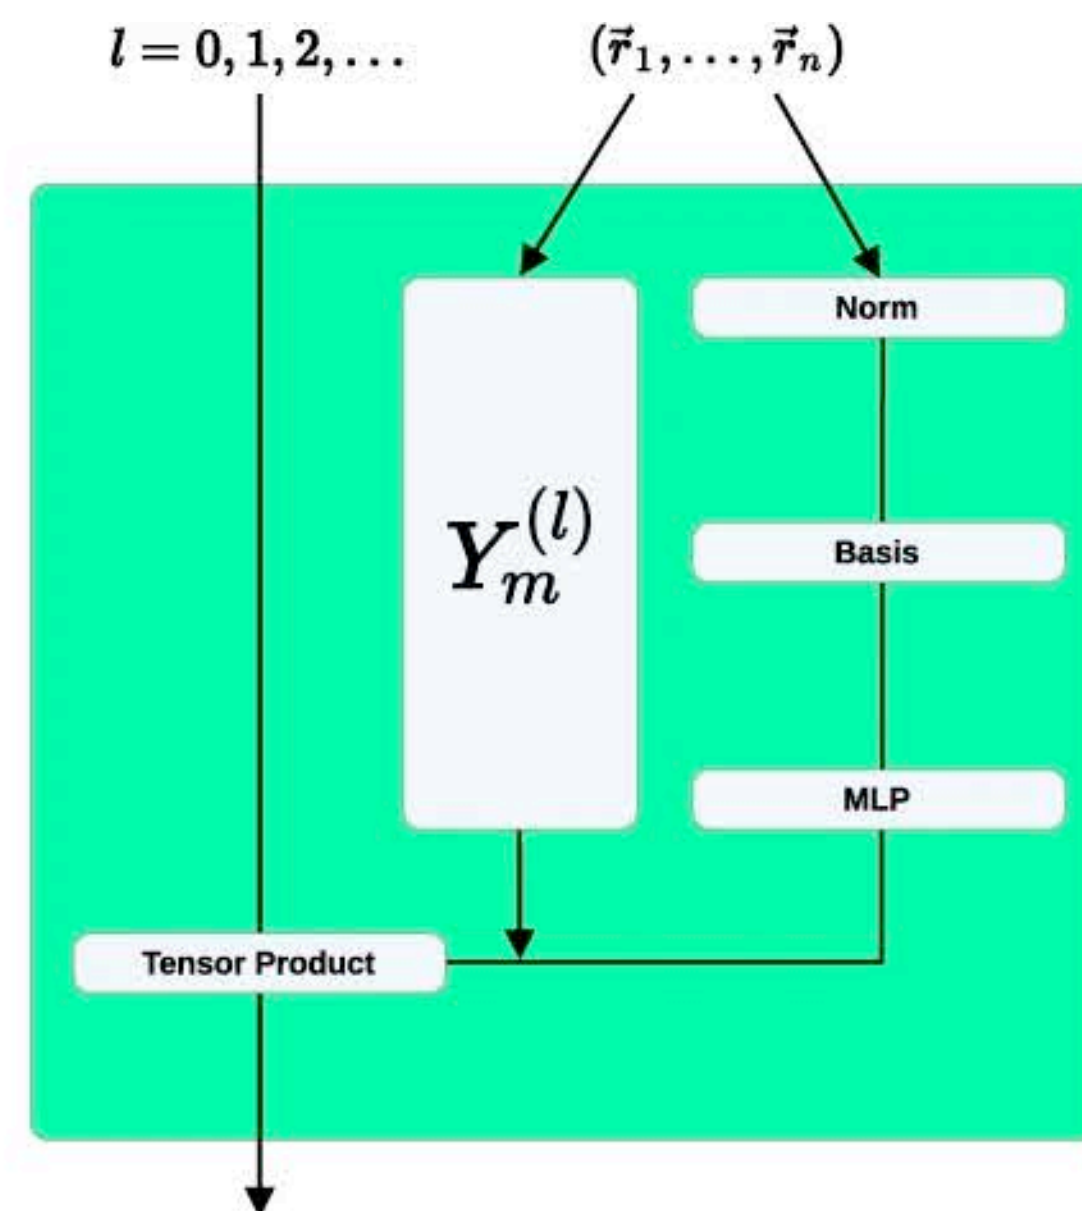
Cormorant



B. Anderson et al. *arXiv:1906.04015* (2019)

Covariant “neurons” for
SO(3) symmetry

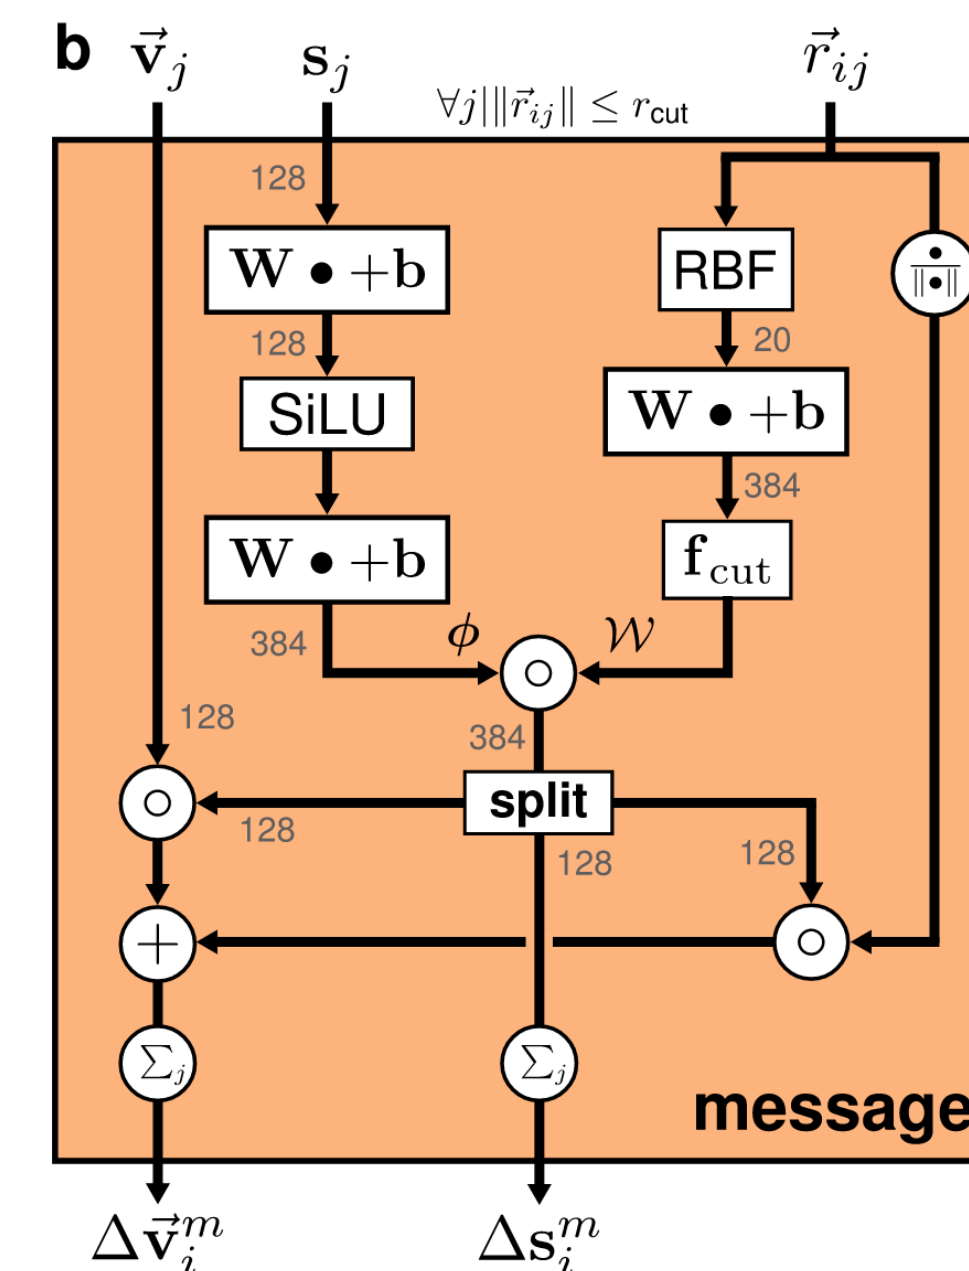
NequIP



S. Batzner et al. *Nat. Commun.* **13**, 2453 (2022)

E(3)-equivariant GNN

PaiNN



K. Schütt et al. *ICML* **139** (2021)

Uses directional message-
passing and vector repr.

These models are close to the state-of-the-art for several datasets

Other improvements: many-body terms and better scaling

Allegro

$$\begin{aligned}
\mathbf{V}_{n,(\ell_1,p_1,\ell_2,p_2)\rightarrow(\ell_{\text{out}},p_{\text{out}})}^{ij,L} &= \sum_{k \in \mathcal{N}(i)} w_{n,\ell_2,p_2}^{ik,L} \left(\mathbf{V}_{n,\ell_1,p_1}^{ij,L-1} \otimes \vec{Y}_{\ell_2,p_2}^{ik} \right) \\
&= \sum_{k \in \mathcal{N}(i)} \mathbf{V}_{n,\ell_1,p_1}^{ij,L-1} \otimes \left(w_{n,\ell_2,p_2}^{ik,L} \vec{Y}_{\ell_2,p_2}^{ik} \right) \\
&= \mathbf{V}_{n,\ell_1,p_1}^{ij,L-1} \otimes \left(\sum_{k \in \mathcal{N}(i)} w_{n,\ell_2,p_2}^{ik,L} \vec{Y}_{\ell_2,p_2}^{ik} \right)
\end{aligned}$$

A. Musaelian et al. *arXiv:2204.05249* (2022)**MACE**


$$\mathbf{m}_i^{(t)} = \sum_j \mathbf{u}_1 \left(\sigma_i^{(t)}; \sigma_j^{(t)} \right) + \sum_{j_1, j_2} \mathbf{u}_2 \left(\sigma_i^{(t)}; \sigma_{j_1}^{(t)}, \sigma_{j_2}^{(t)} \right) + \dots + \sum_{j_1, \dots, j_\nu} \mathbf{u}_\nu \left(\sigma_i^{(t)}; \sigma_{j_1}^{(t)}, \dots, \sigma_{j_\nu}^{(t)} \right),$$

$$A_{i,kl_1 m_1}^{(1)} = \sum_{j \in \mathcal{N}(i)} R_{kl_1}^{(1)}(r_{ji}) Y_{l_1}^{m_1}(\hat{\mathbf{r}}_{ji}) \mathcal{W}_{kz_j}^{(1)}.$$

I. Batatia et al. *arXiv:2206.07697* (2022)

The math gets complicated, but the models get more accurate (and scale better)

To summarize what we have learned so far

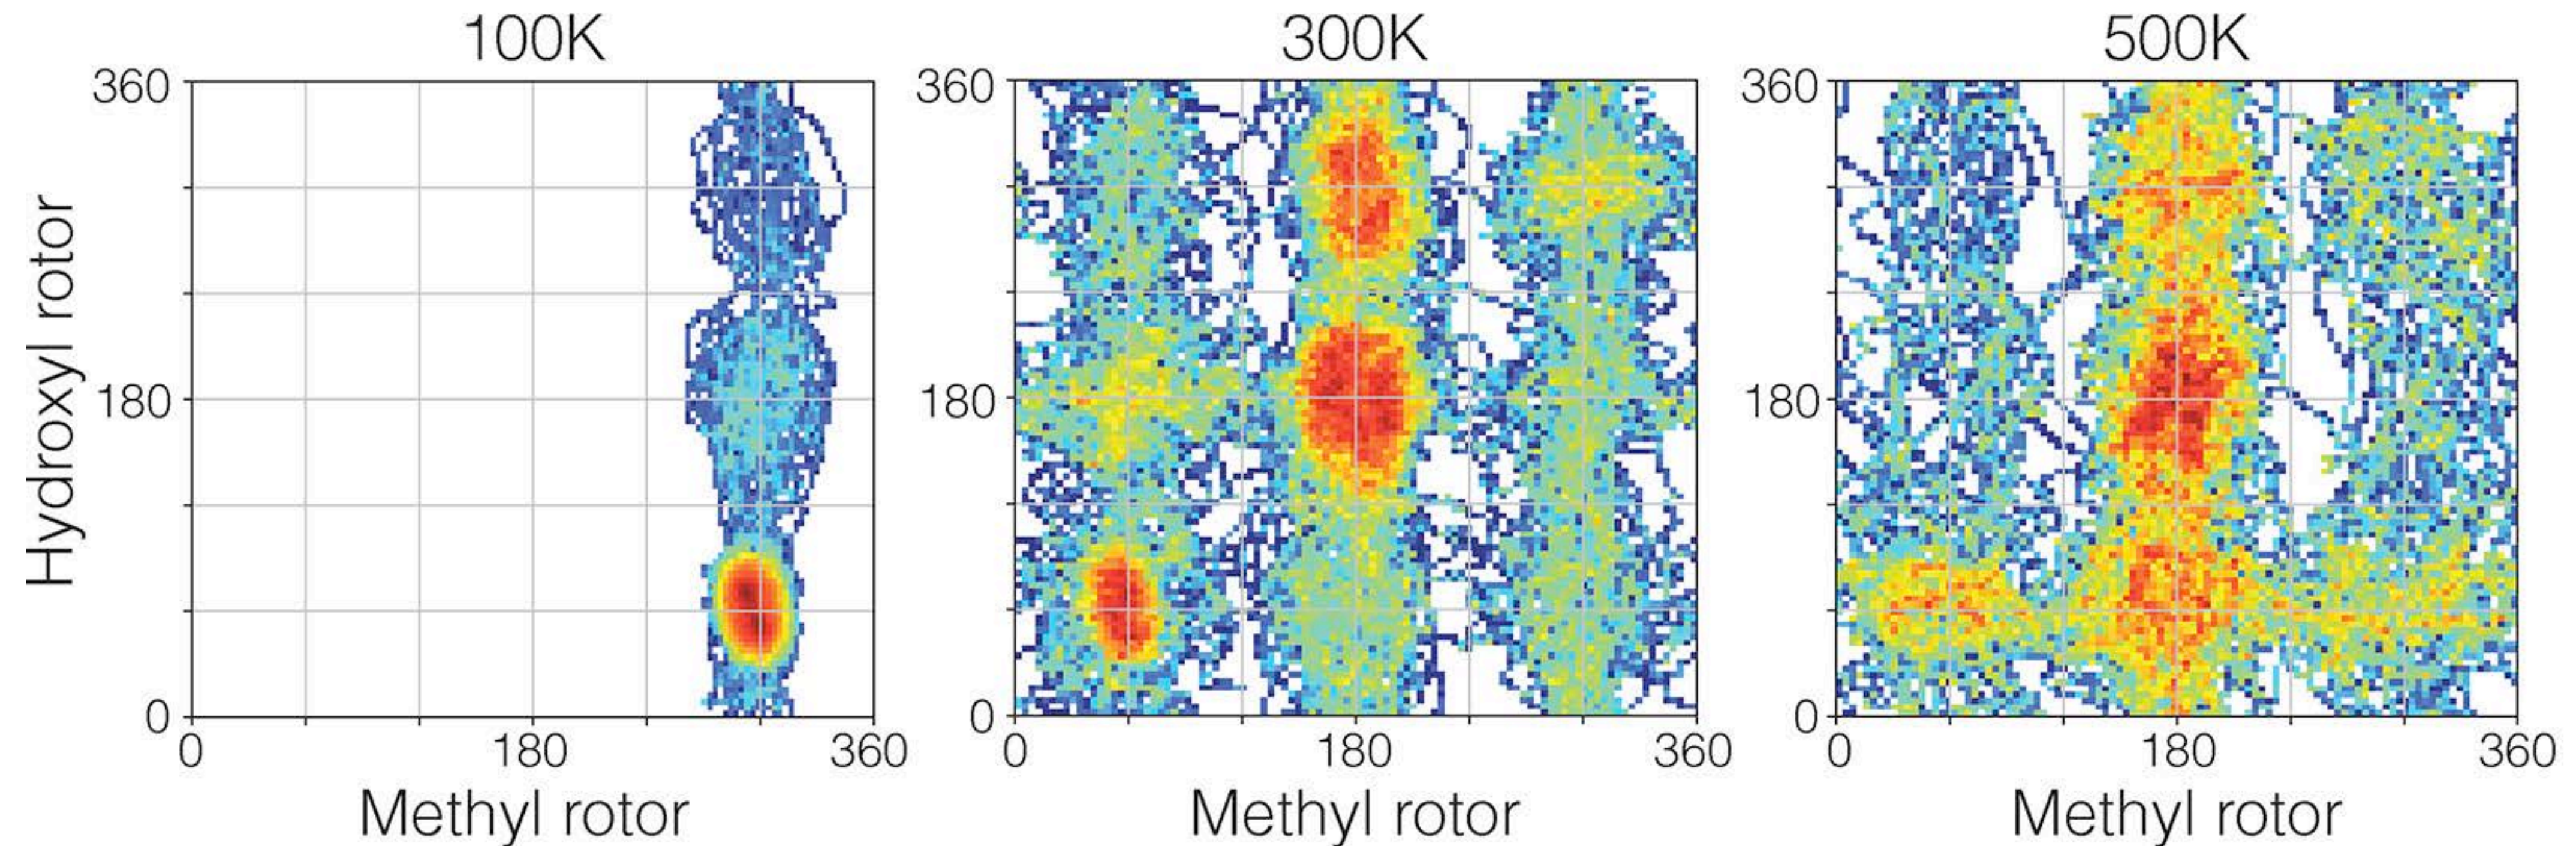


4. Data

How to use, construct, and validate datasets in NNIPs?

Machine learning needs data

For example, if we wanted to perform an MD simulation for ethanol, we would observe that the sampling of configuration space changes with the temperature, as expected:



O. Unke et al. *Chem. Rev.* **121** (16), 10142 (2021)

To train a NN force field, we have to use the right datasets for our application of interest.

How to create data?

Create data as you usually would: QM, DFT, etc.

AIMD

: use AIMD trajectories as dataset

: easy to perform

: high cost, correlated samples

normal mode sampling

: displace atoms randomly along the eigenvectors of the Hessian


: easy to perform

: small distortions only

enhanced sampling

: explore the PES with enhanced sampling methods (e.g., metadynamics)

: better exploration of the PES

: harder to implement, often relying on *ab initio*

active learning

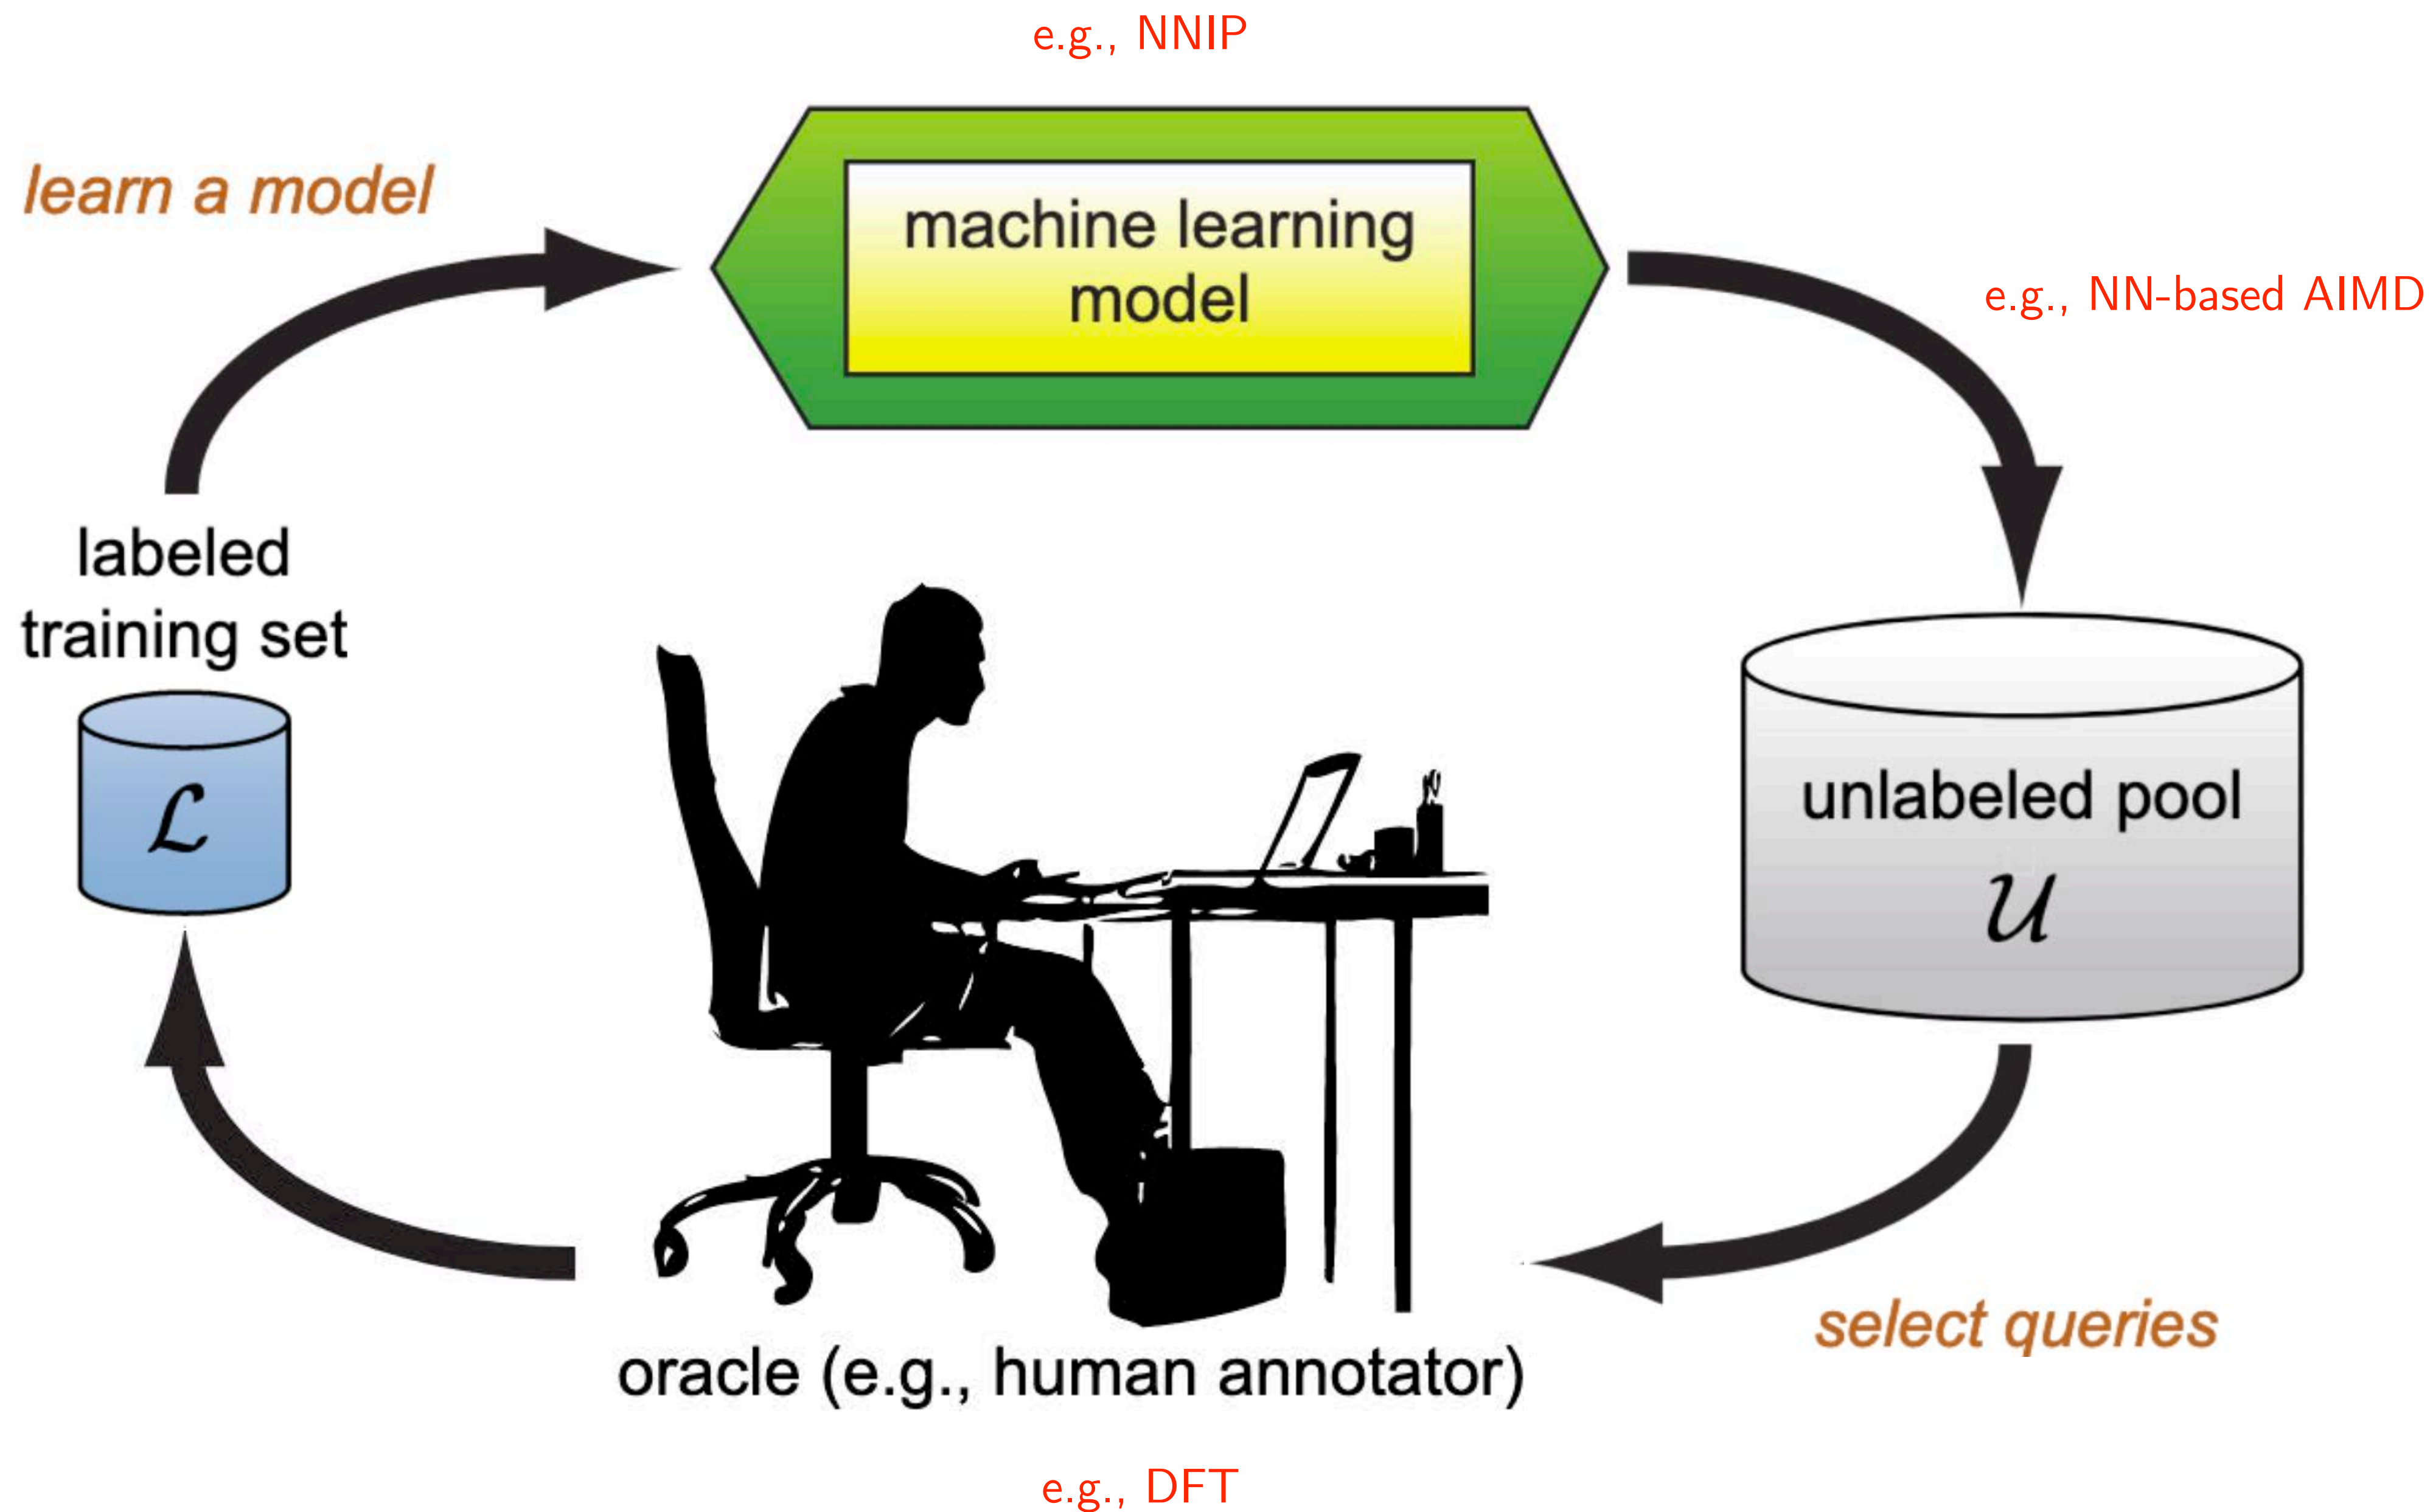
: improve the dataset over time by analyzing the uncertainties

: good quality datasets, may be cheaper to produce

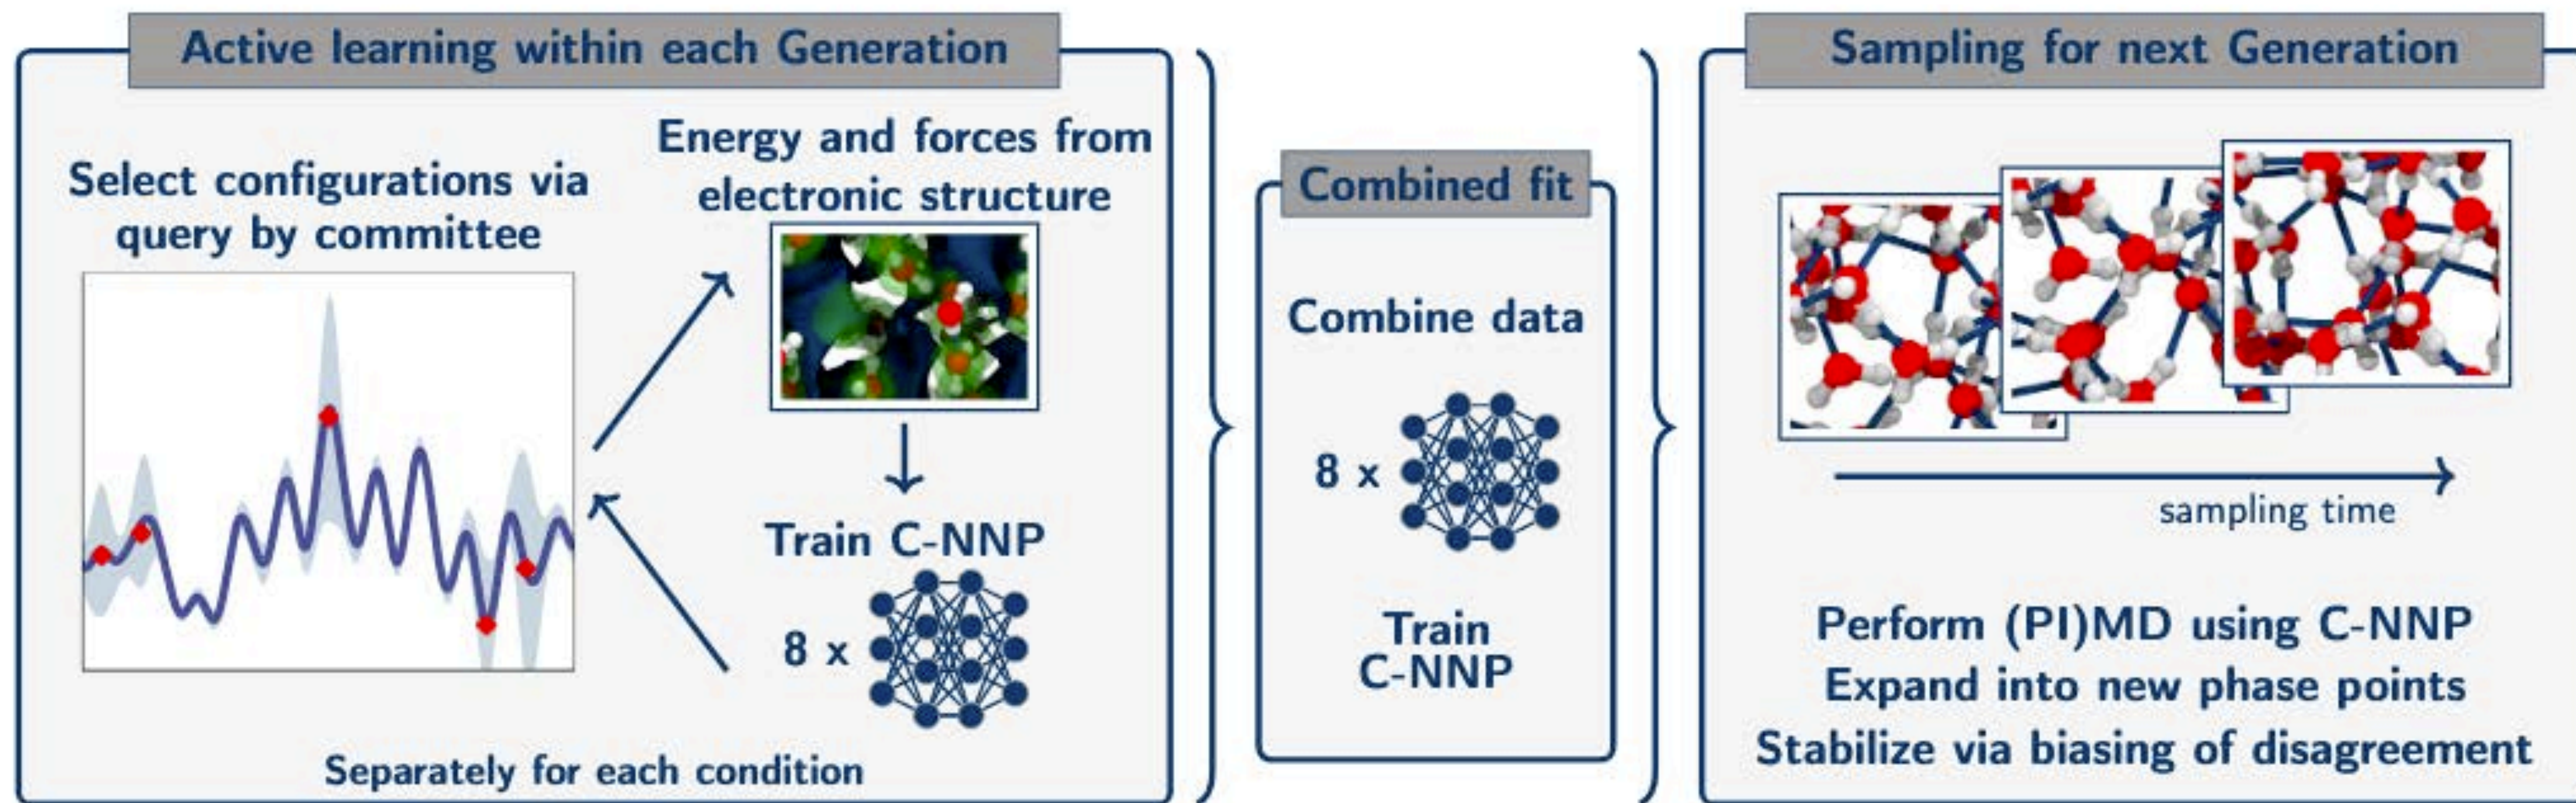
: requires uncertainty metric, long process

let's focus on this one

What is active learning?

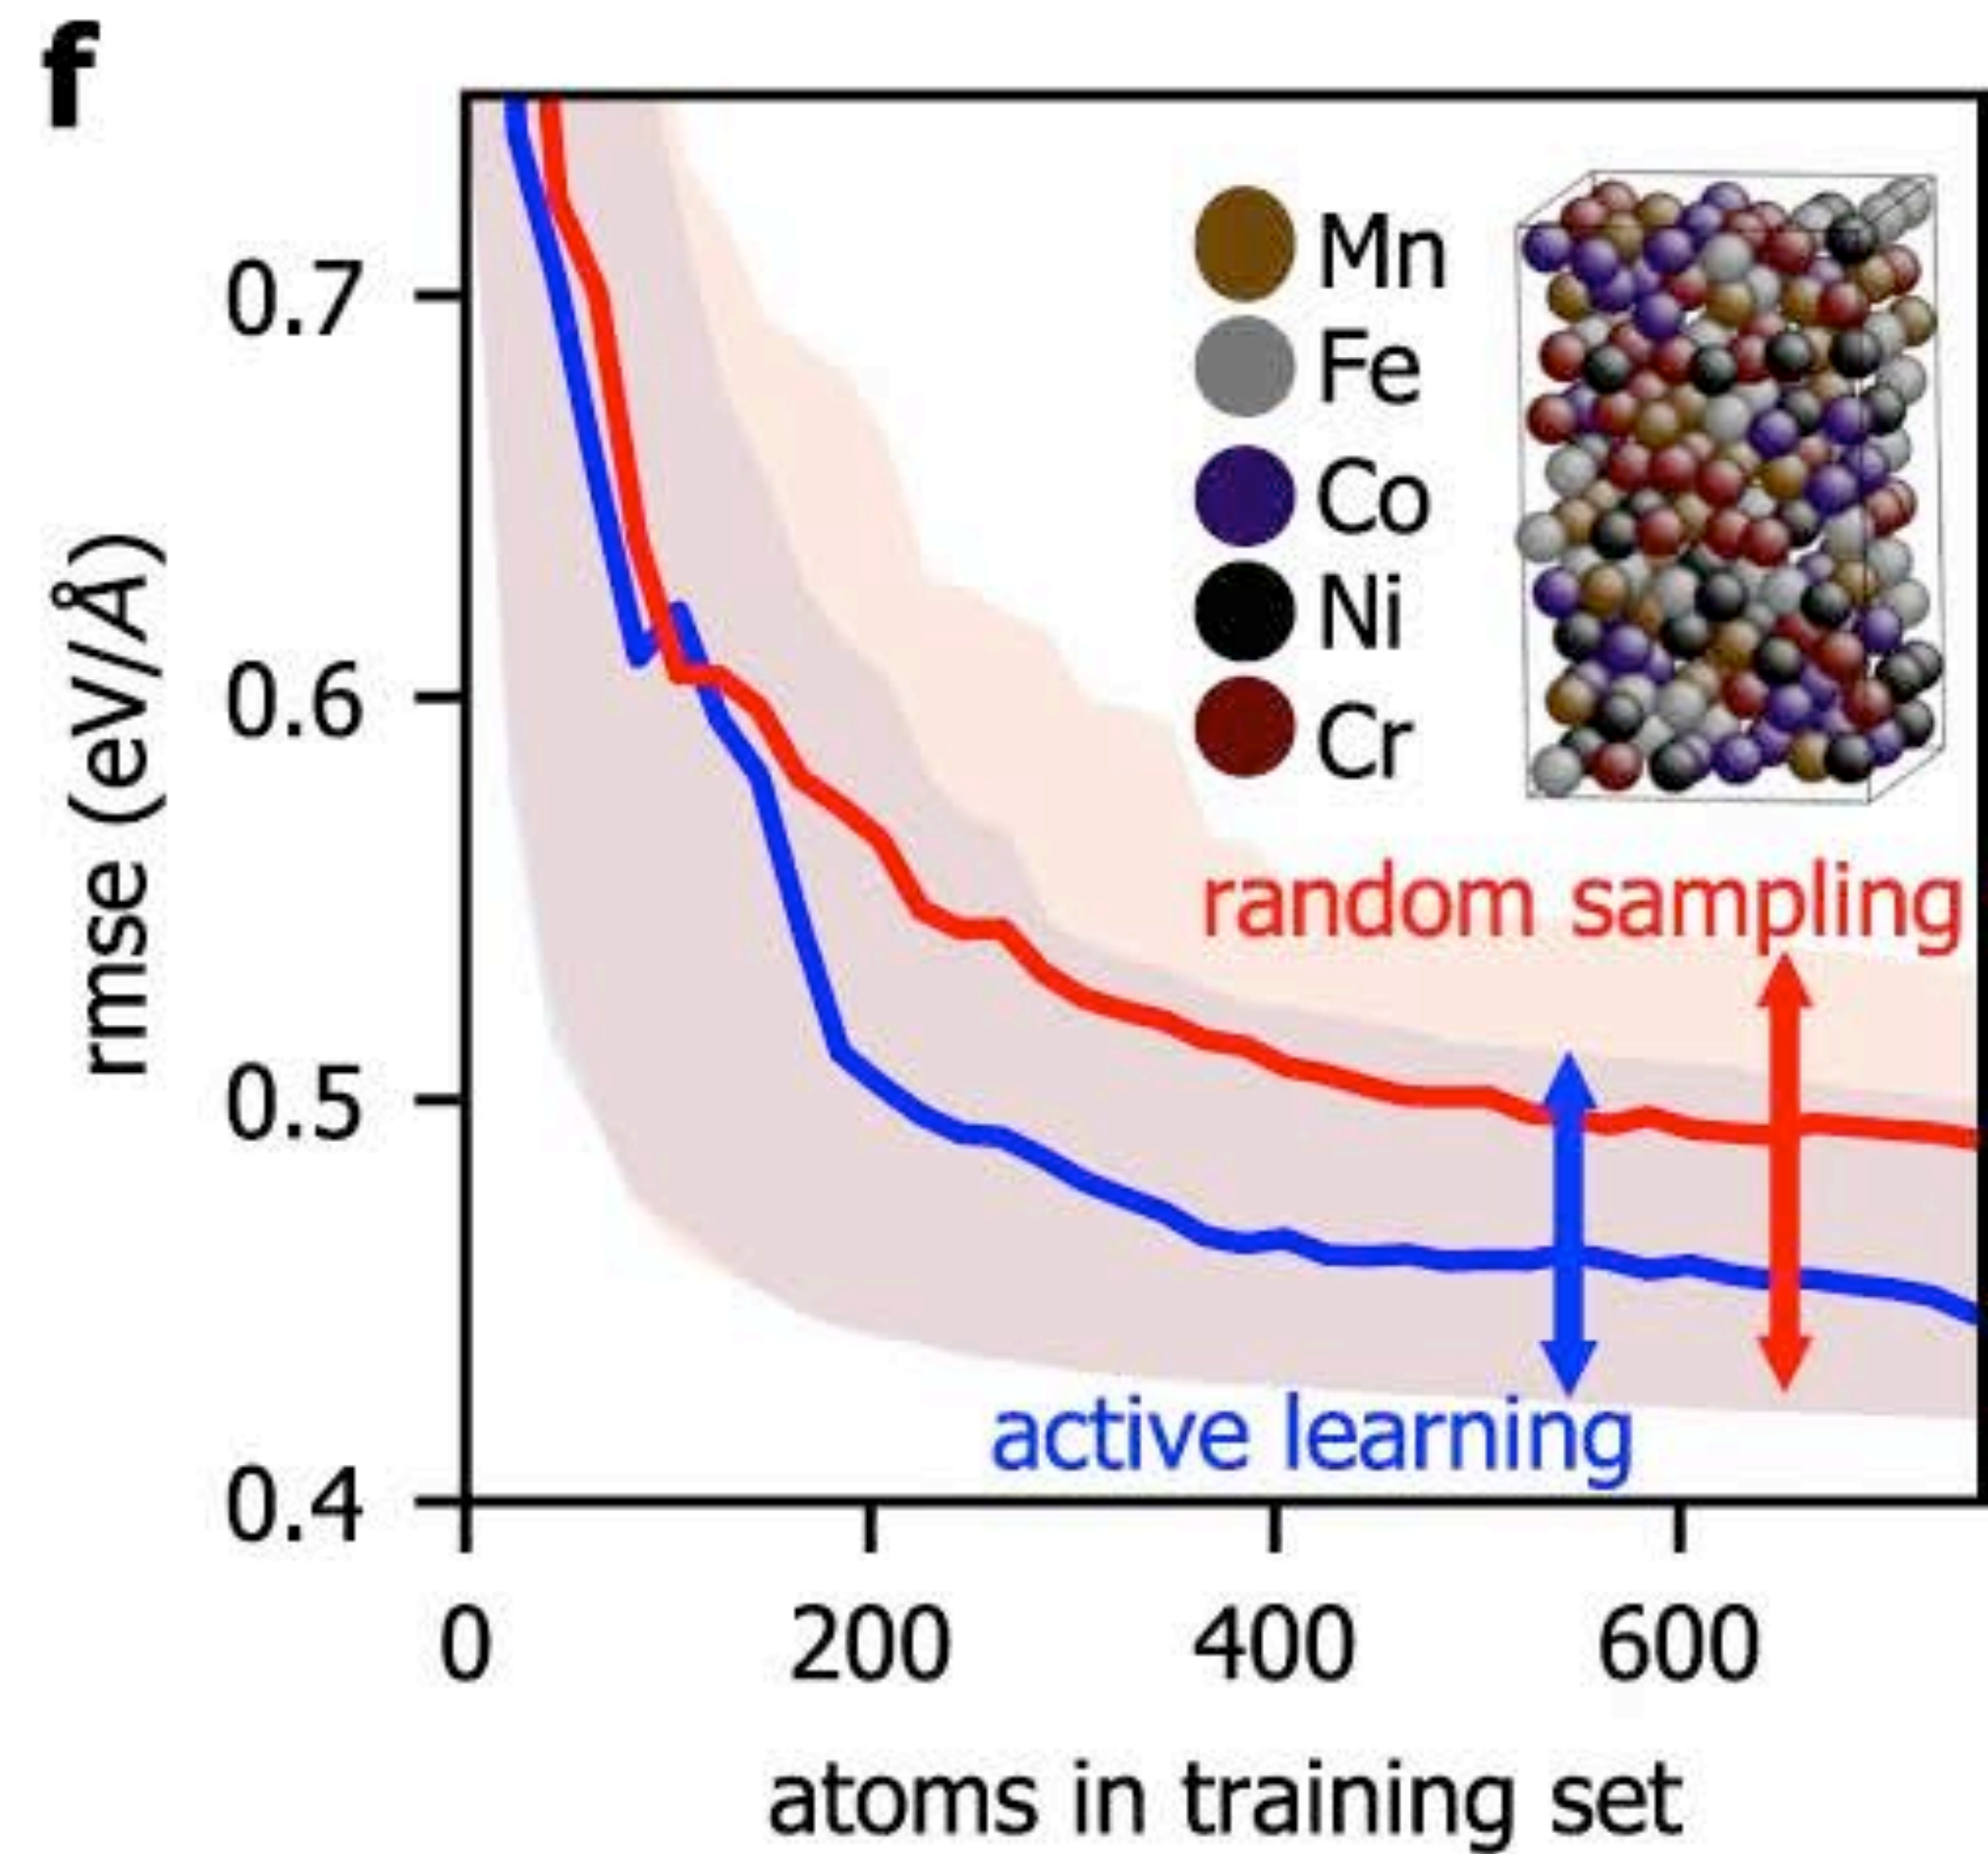
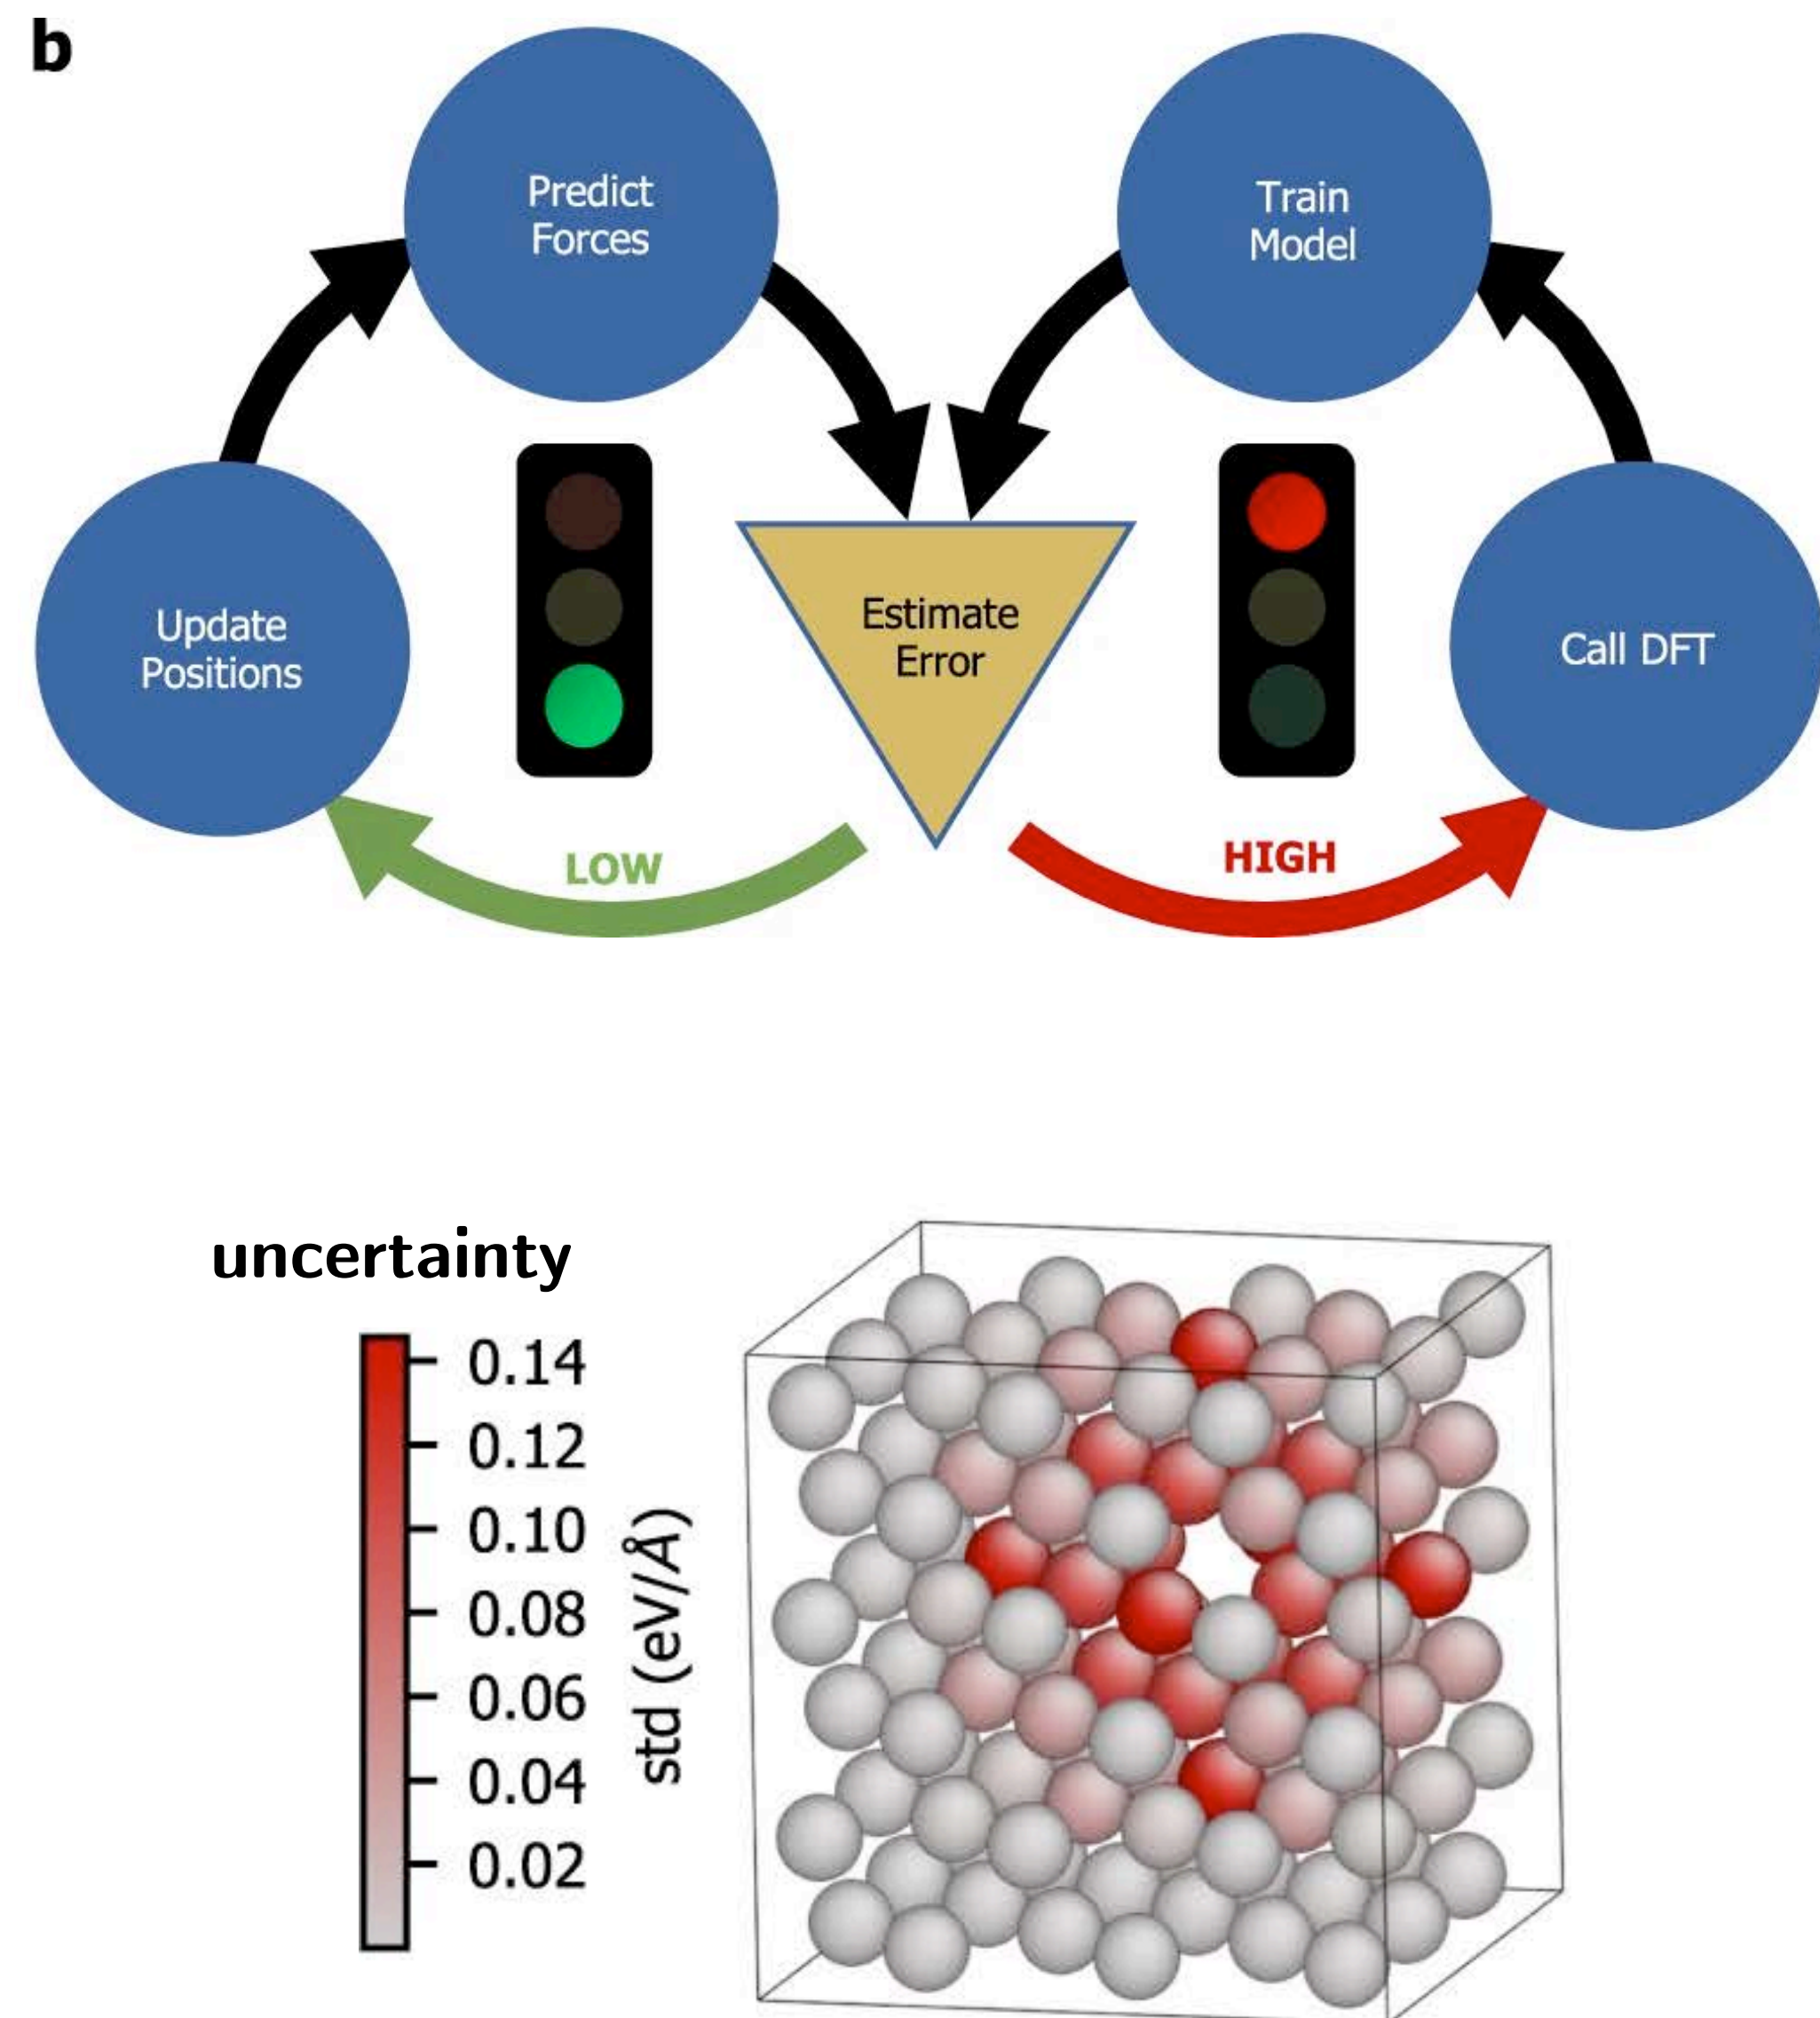


Active learning requires uncertainty quantification to identify “unlabeled” configurations

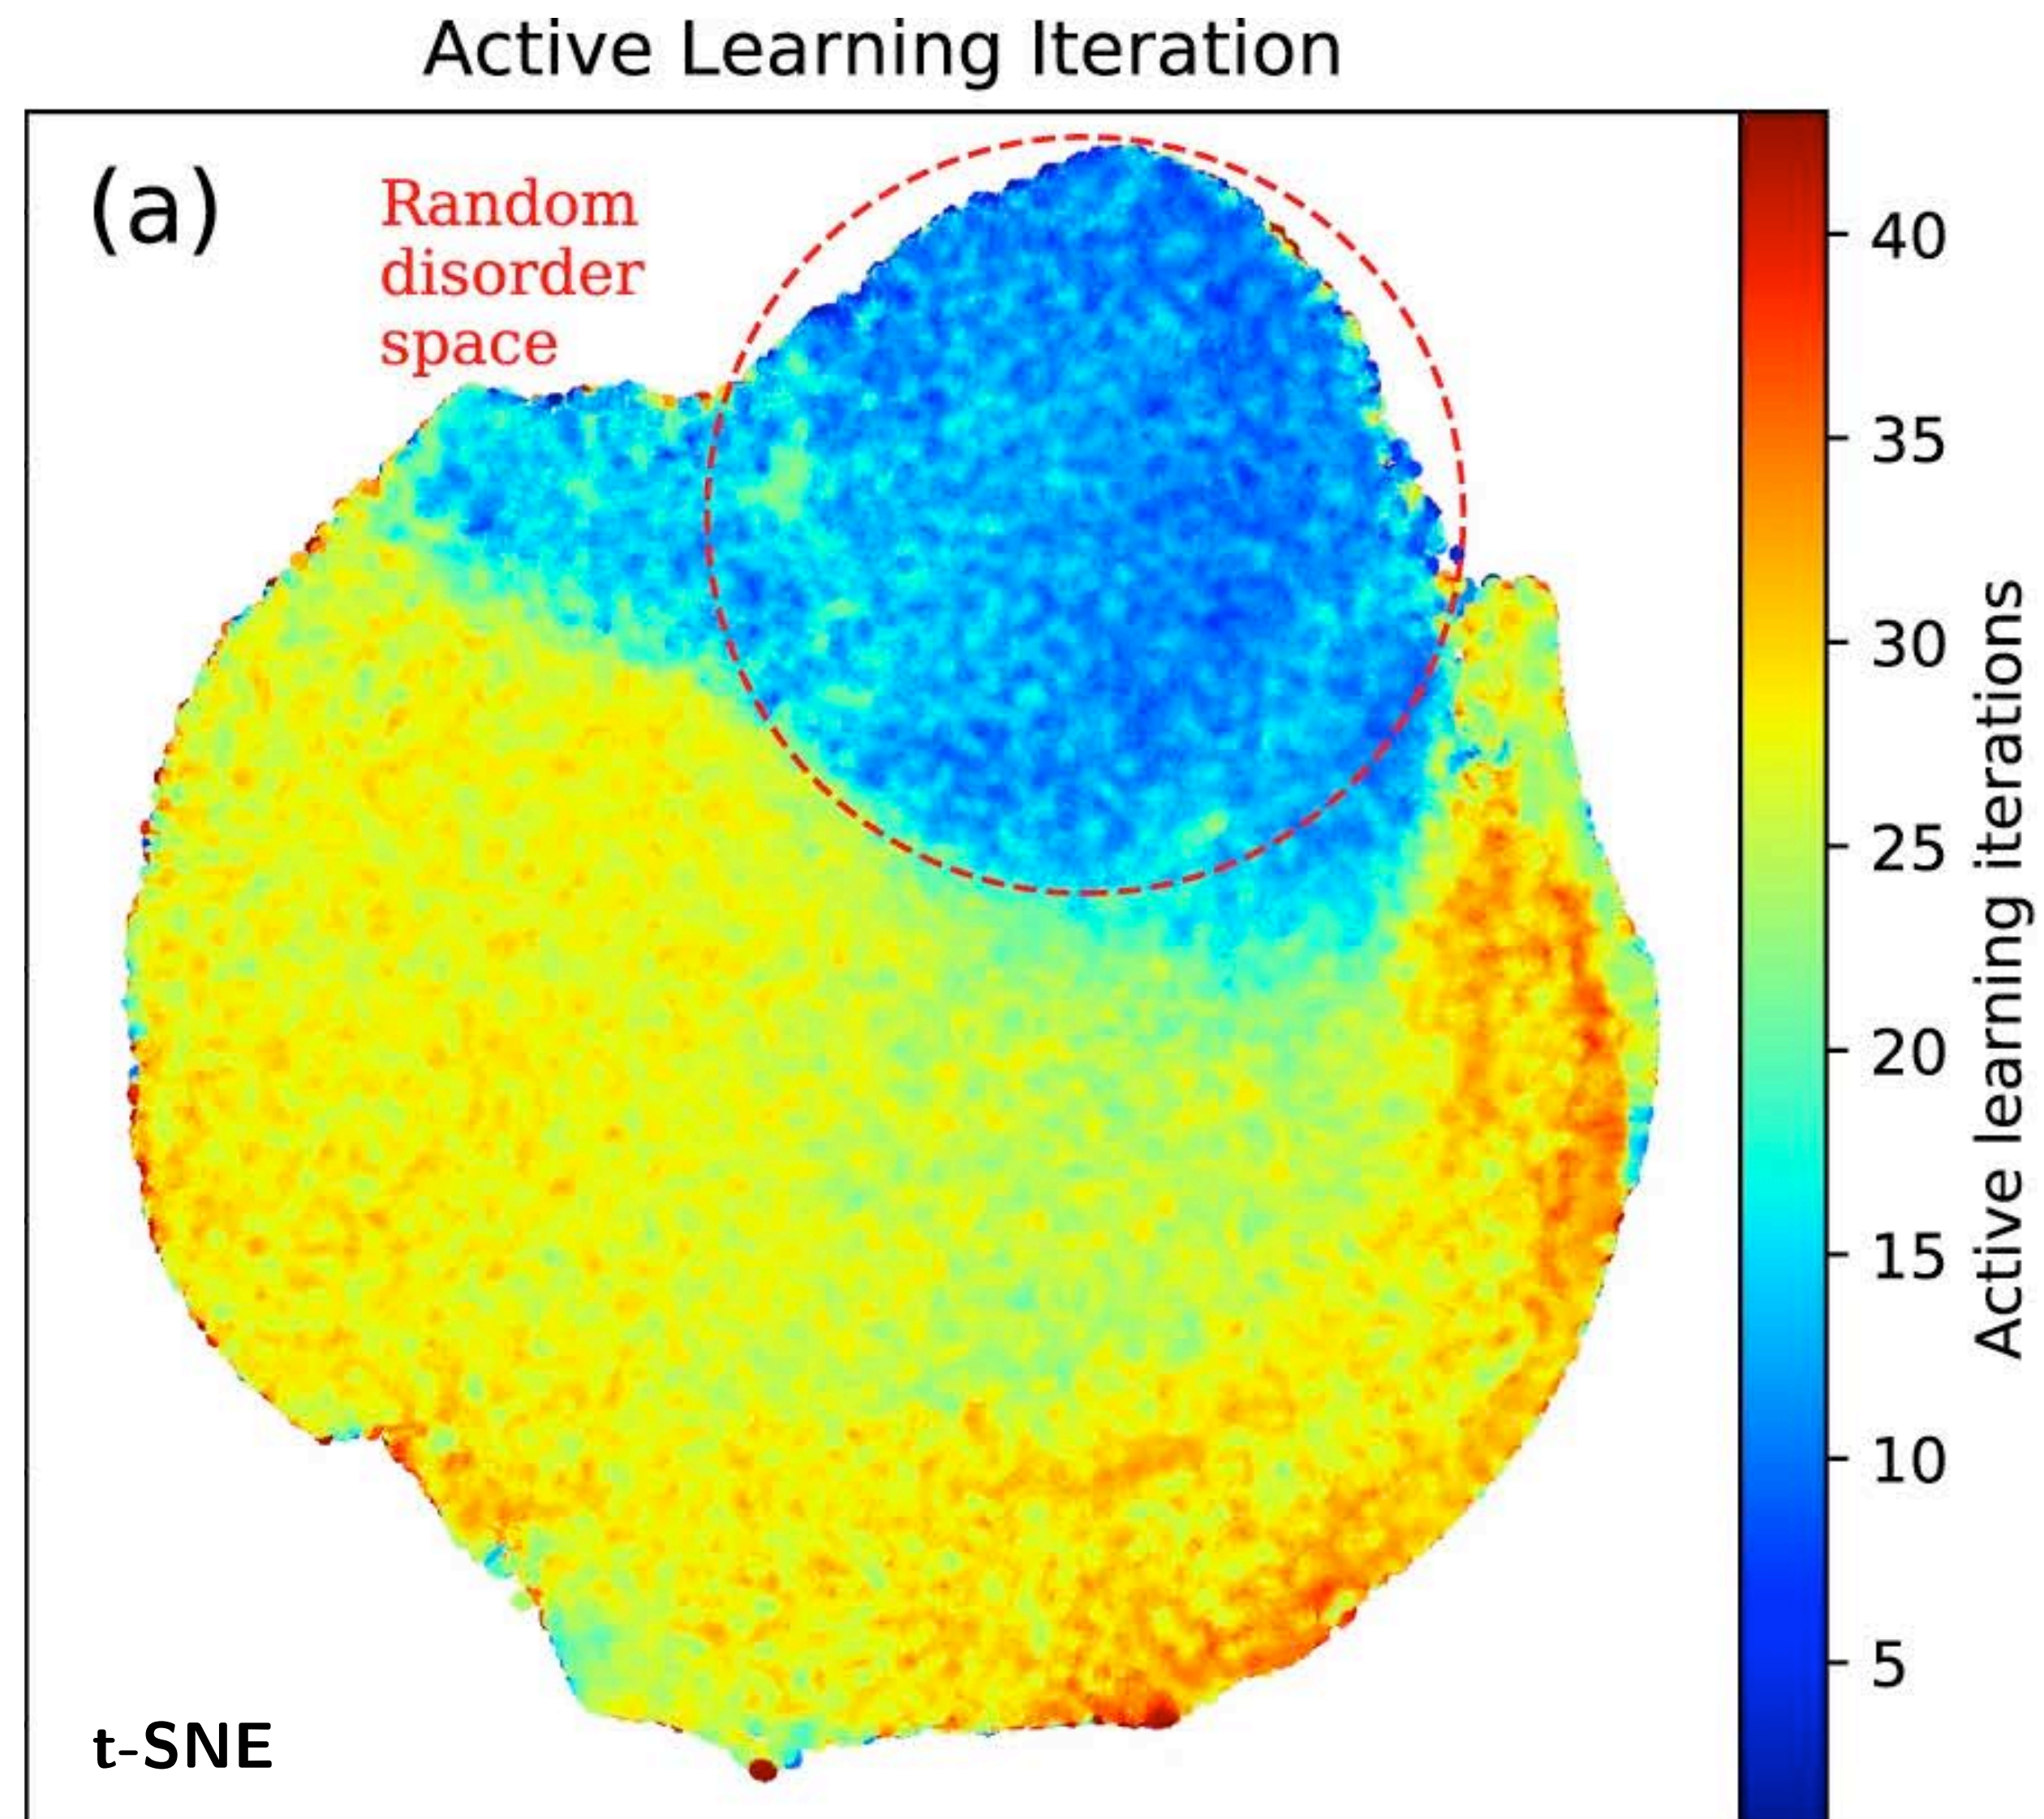
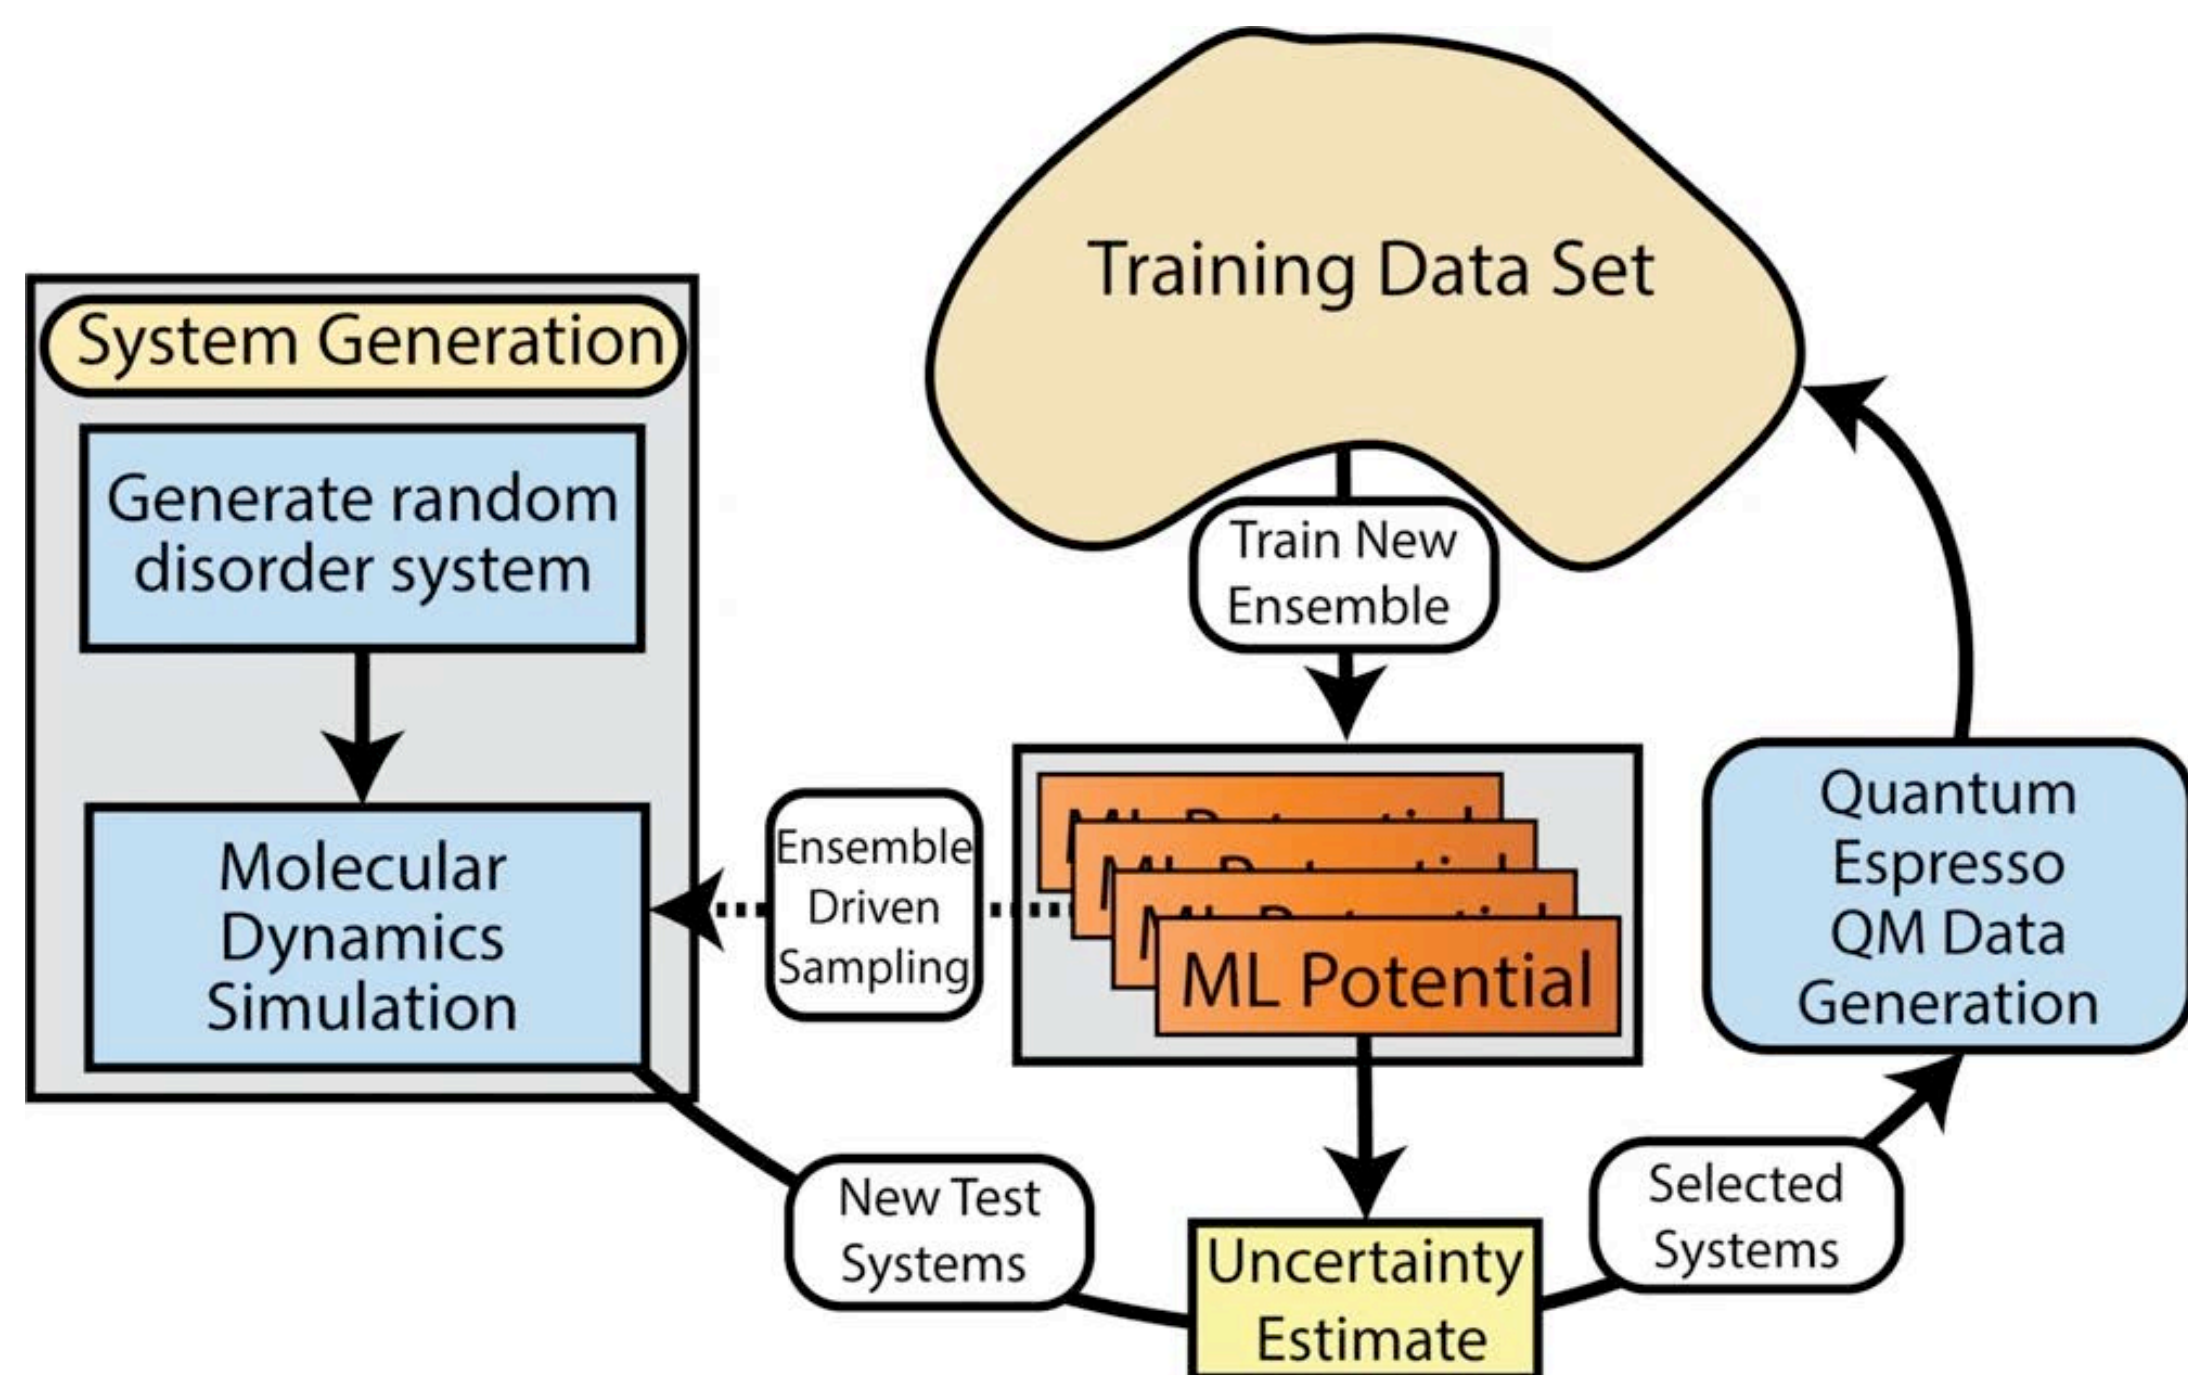


What is the problem of this?

Examples of interesting dataset constructions: GPR, FLARE and high-entropy alloys




Examples of interesting dataset constructions: aluminum dataset from ANI potential



Some datasets typically used out there


rMD17



S. Chmiela et al. *Sci. Adv.* **3** (5), e1603015 (2017)
A. Christiansen et al. *MLST* **1**, 045018 (2020)

Small molecules with
their conformers


3BPA



D. Kovács et al. *JCTC* **17** (12), 7696 (2021)

Flexible molecule sampled
at different temperatures

Open Catalyst (OC20)




L. Chanussot et al. *ACS Catalysis* **11**, 6059 (2021)

Adsorption energies on
inorganic catalysts

Remember: train-validation-test, k-fold CV, error metrics and many more.

There are many, many datasets...

To summarize what we have learned so far

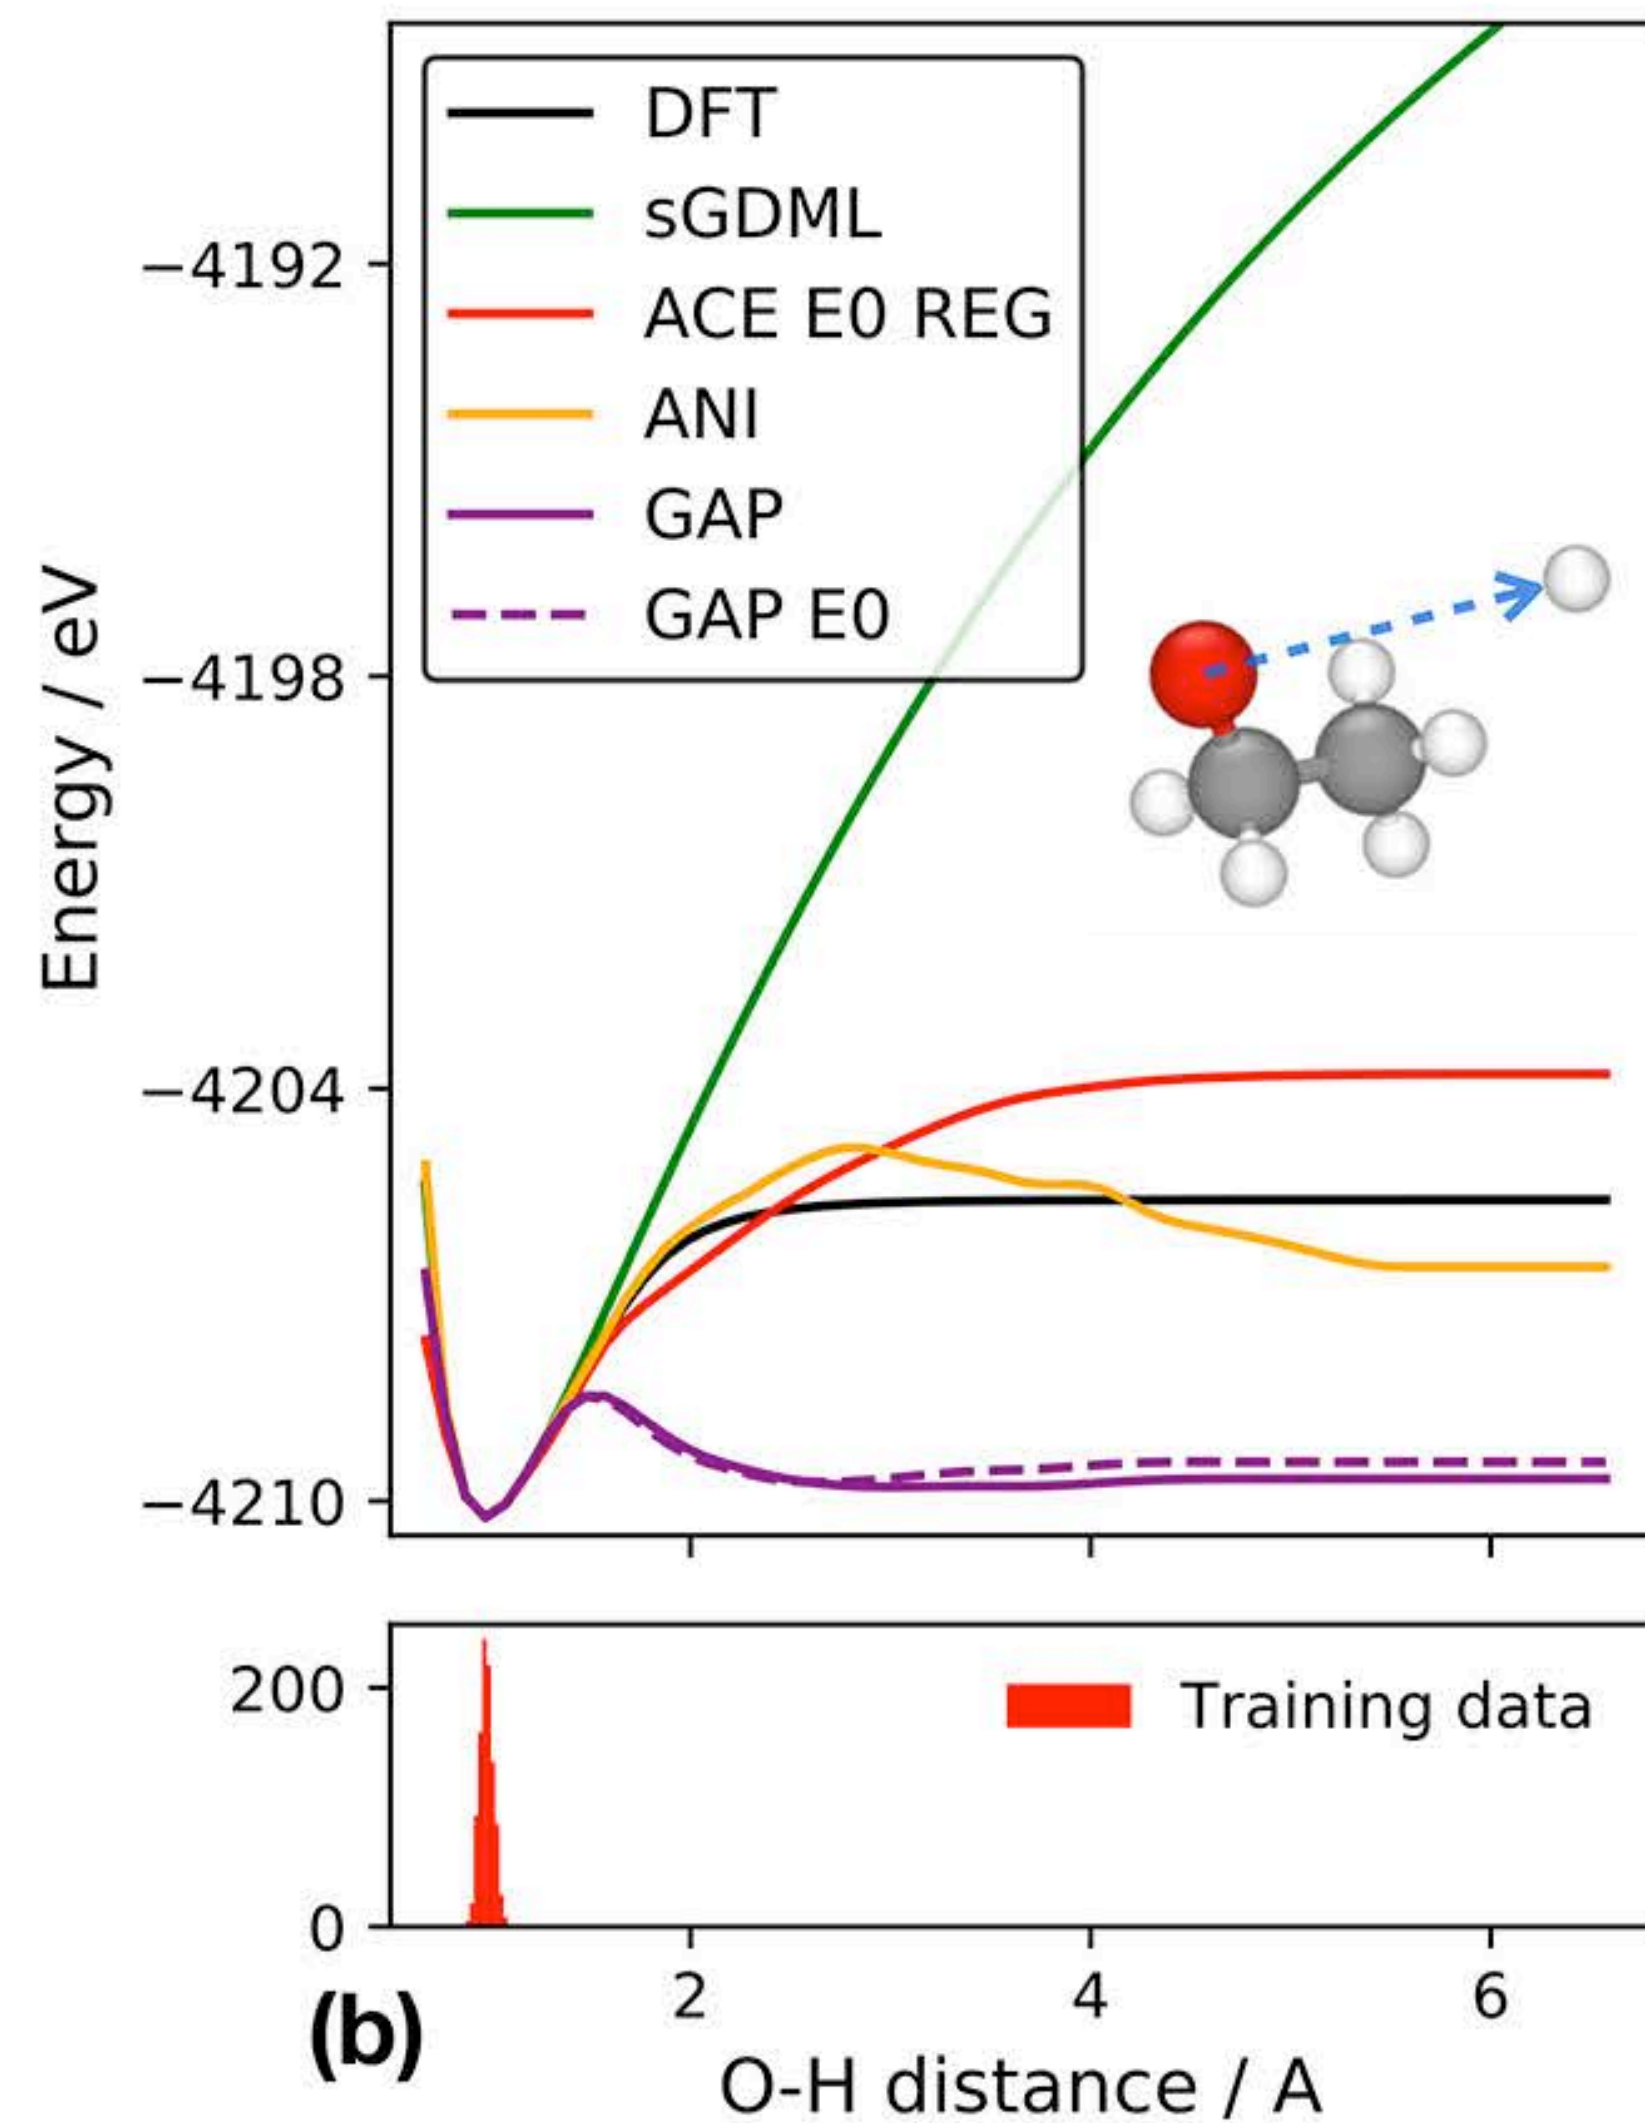


5. Frontiers of NNIPs


What's next for ML interatomic potentials?

Interesting trends in the field


Uncertainty & extrapolation

D. Kovács et al. *JCTC* **17** (12), 7696 (2021)

Challenging applications


J. Westermayr et al. *Nat. Chem.* **14**, 914 (2022)

Better metrics & interpretability



J. Vita and D. Schwalbe-Koda (2023)

For example: ML potentials suffer when extrapolating...






x

“panda”

57.7% confidence

+ .007 ×



$\text{sign}(\nabla_x J(\theta, x, y))$

“nematode”

8.2% confidence

=



$x +$

$\epsilon \text{sign}(\nabla_x J(\theta, x, y))$

“gibbon”

99.3 % confidence

$$\min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\max_{\delta \in \Delta} \mathcal{L}(h_{\theta}(x + \delta), y) \right]$$

find the NN weights that minimize

across the whole dataset under study


and for perturbations δ in the set of allowed perturbations Δ

the perturbed loss function

The diagram illustrates the optimization process for finding robust NN weights. It features a central equation: $\min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\max_{\delta \in \Delta} \mathcal{L}(h_{\theta}(x + \delta), y) \right]$. Four arrows point upwards from explanatory text to parts of the equation: one from 'find the NN weights that minimize' to the θ parameter; one from 'across the whole dataset under study' to the expectation operator $\mathbb{E}_{(x,y) \sim \mathcal{D}}$; one from 'and for perturbations δ in the set of allowed perturbations Δ ' to the maximization operator $\max_{\delta \in \Delta}$; and one from 'the perturbed loss function' to the loss function $\mathcal{L}(h_{\theta}(x + \delta), y)$.

Qualitative results of robust NNs

Original



primate

Standard



dog

 ℓ_2 -trained

dog



bird



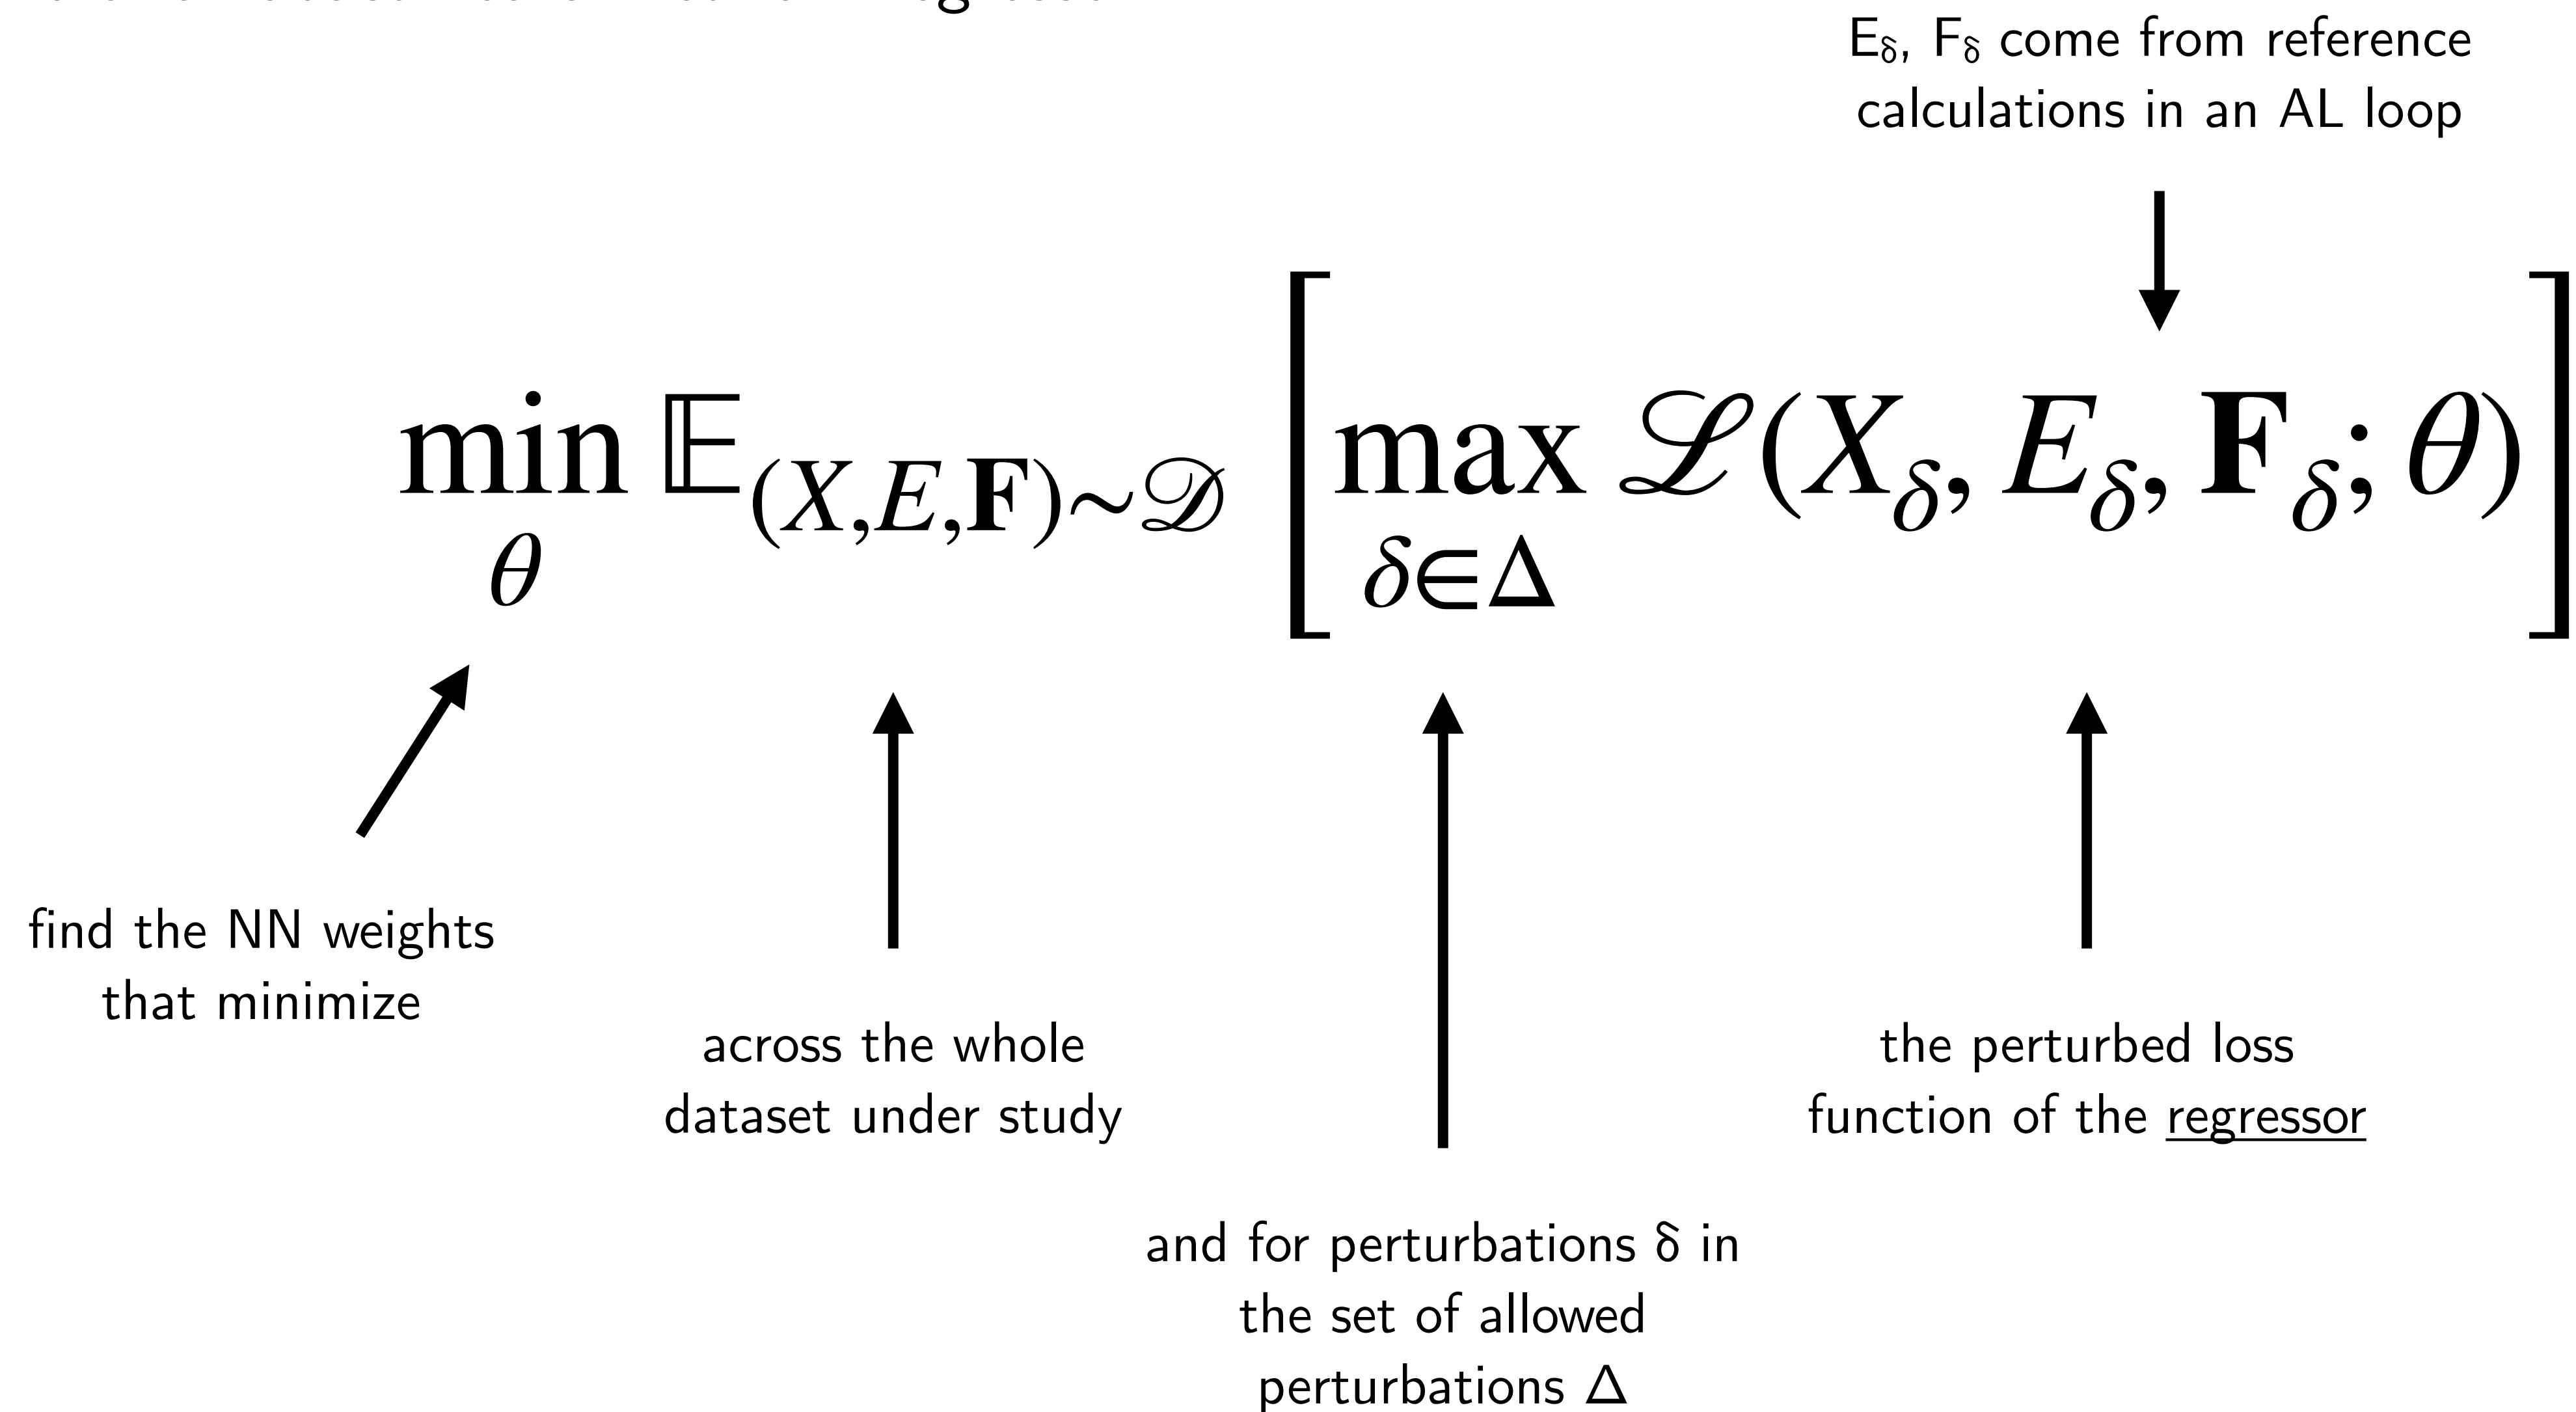
turtle



cat

Question: how to do this for NN potentials?

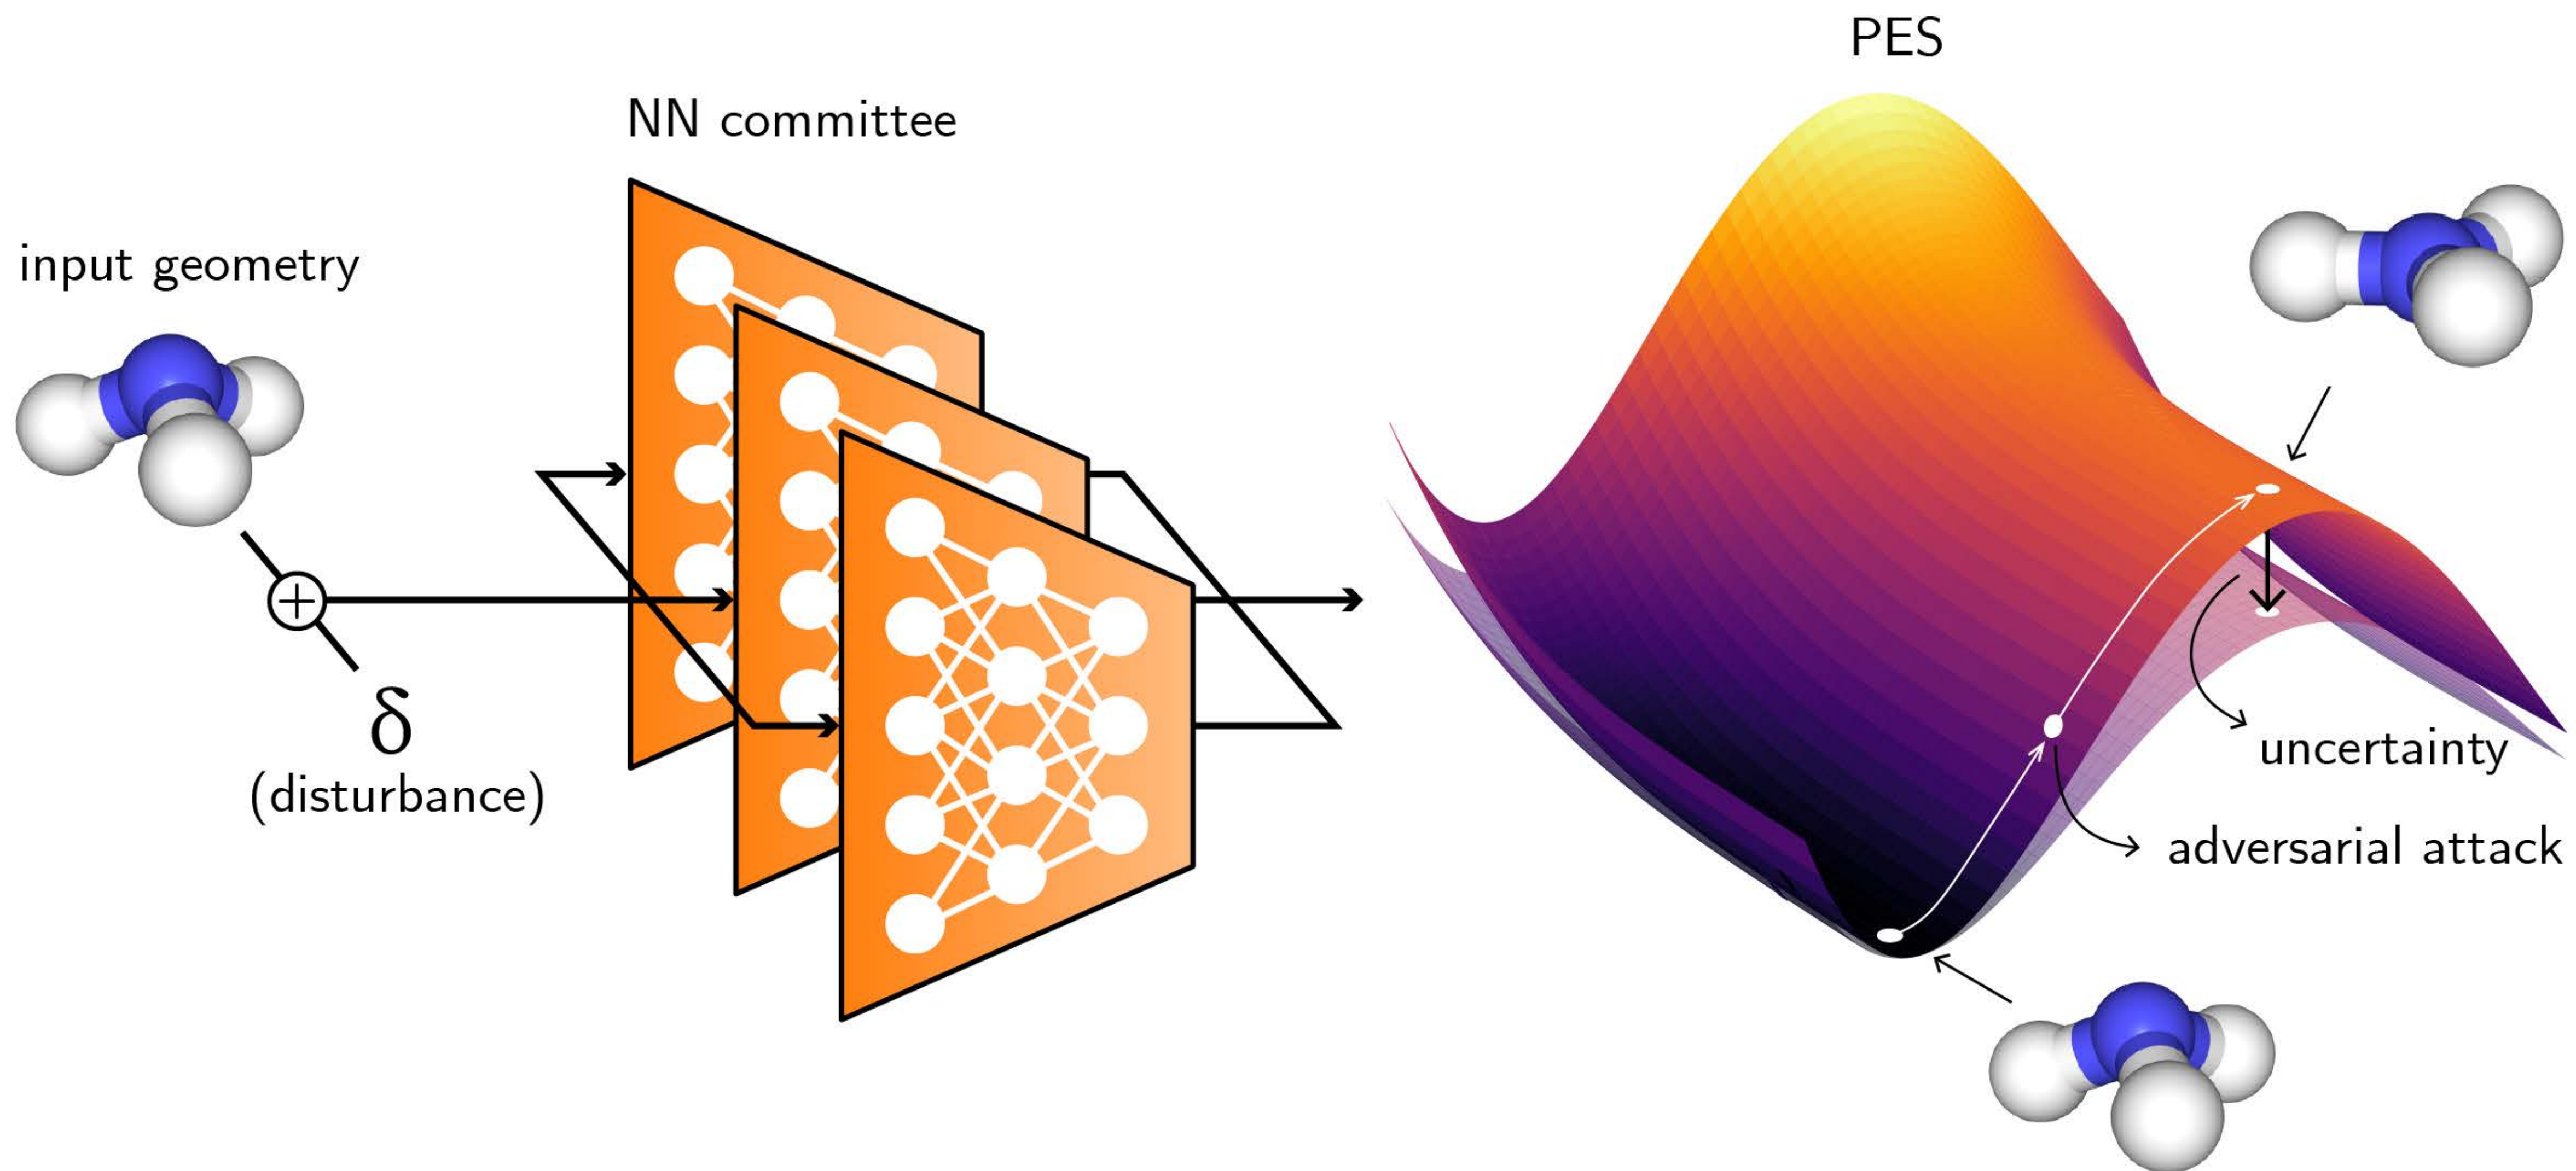
Objective of a robust neural network regressor



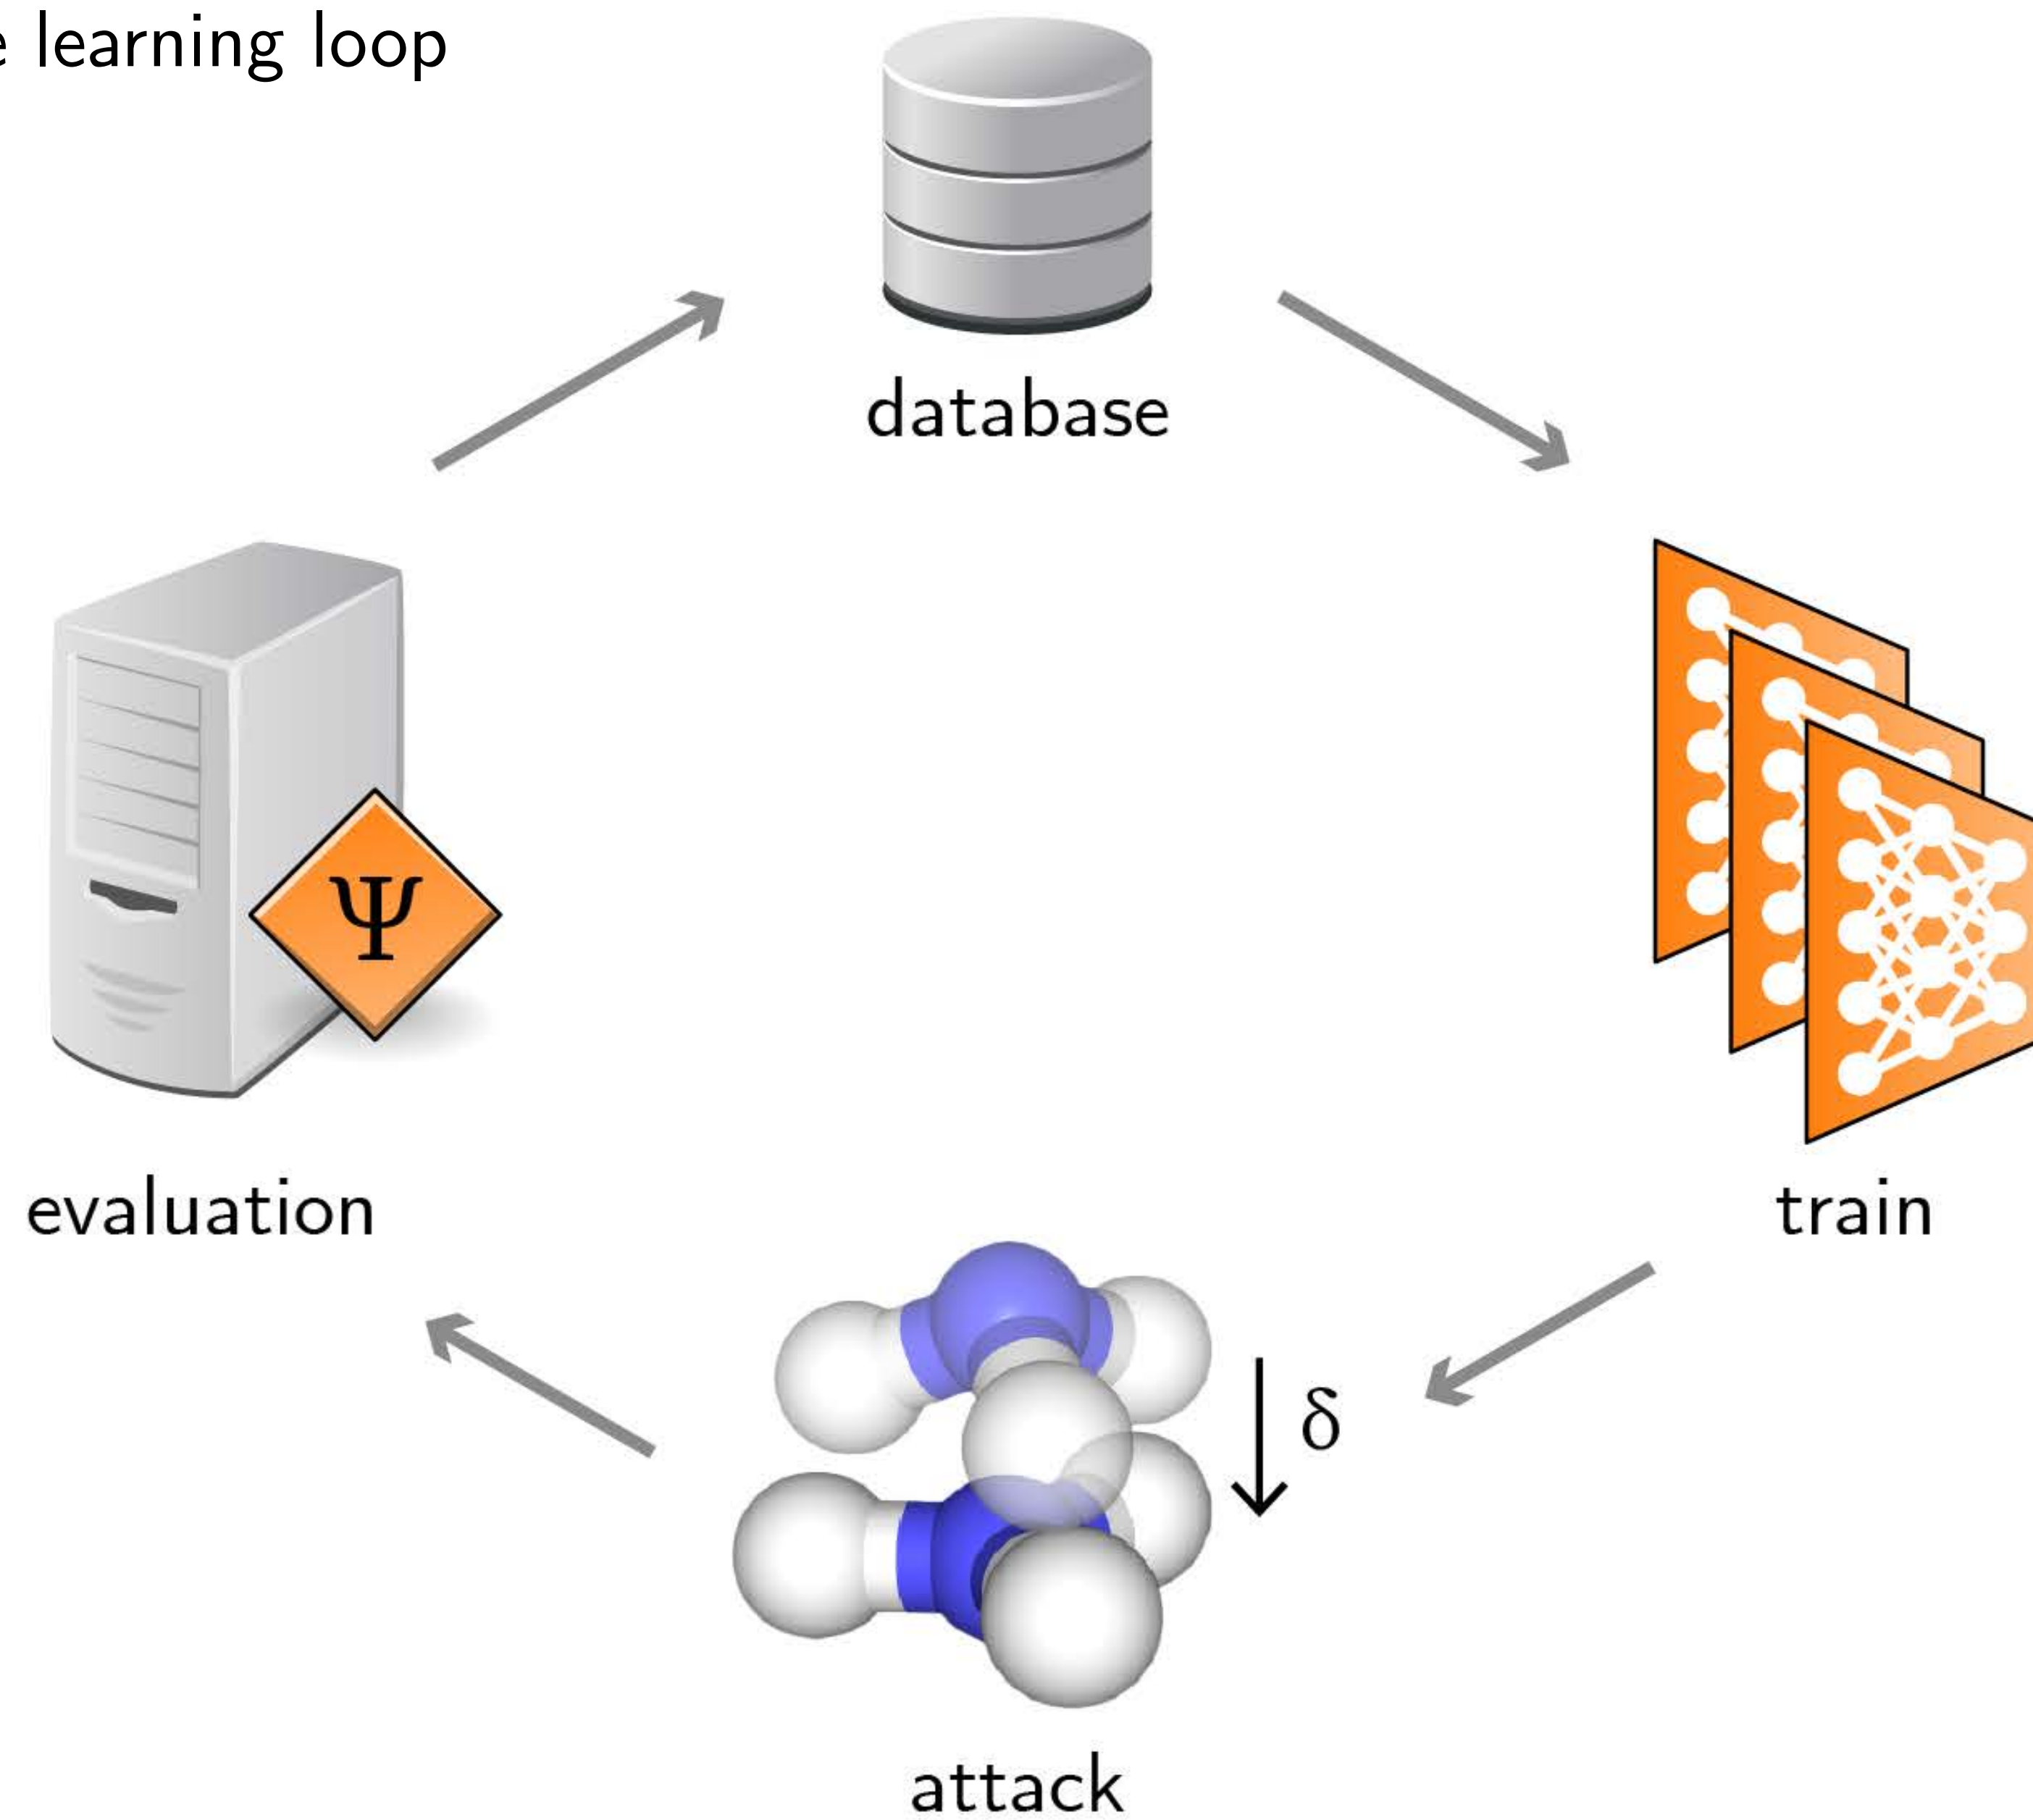
Question: how to generate the perturbed samples and their ground truth values?

Idea: find geometries that maximize the epistemic uncertainty of the NN potential!


Adversarial loss depends on the uncertainty



Robust training is an active learning loop



Sample new points through adversarial attacks



Goals of a good adversarial attack:

- Find points of maximum uncertainty
- Penalize going towards crazy high energies

$$\min_{\theta} \mathbb{E}_{(X, E, \mathbf{F}) \sim \mathcal{D}} \left[\max_{\delta \in \Delta} \mathcal{L}(X_{\delta}, E_{\delta}, \mathbf{F}_{\delta}; \theta) \right]$$

Loss function for adversarial attack

Construct partition function from training set:


$$Q = \sum_{(X, E, \mathbf{F}) \in \mathcal{D}} \exp\left(\frac{-E}{kT}\right)$$

Estimate Boltzmann probability given the mean energy from NN ensemble:

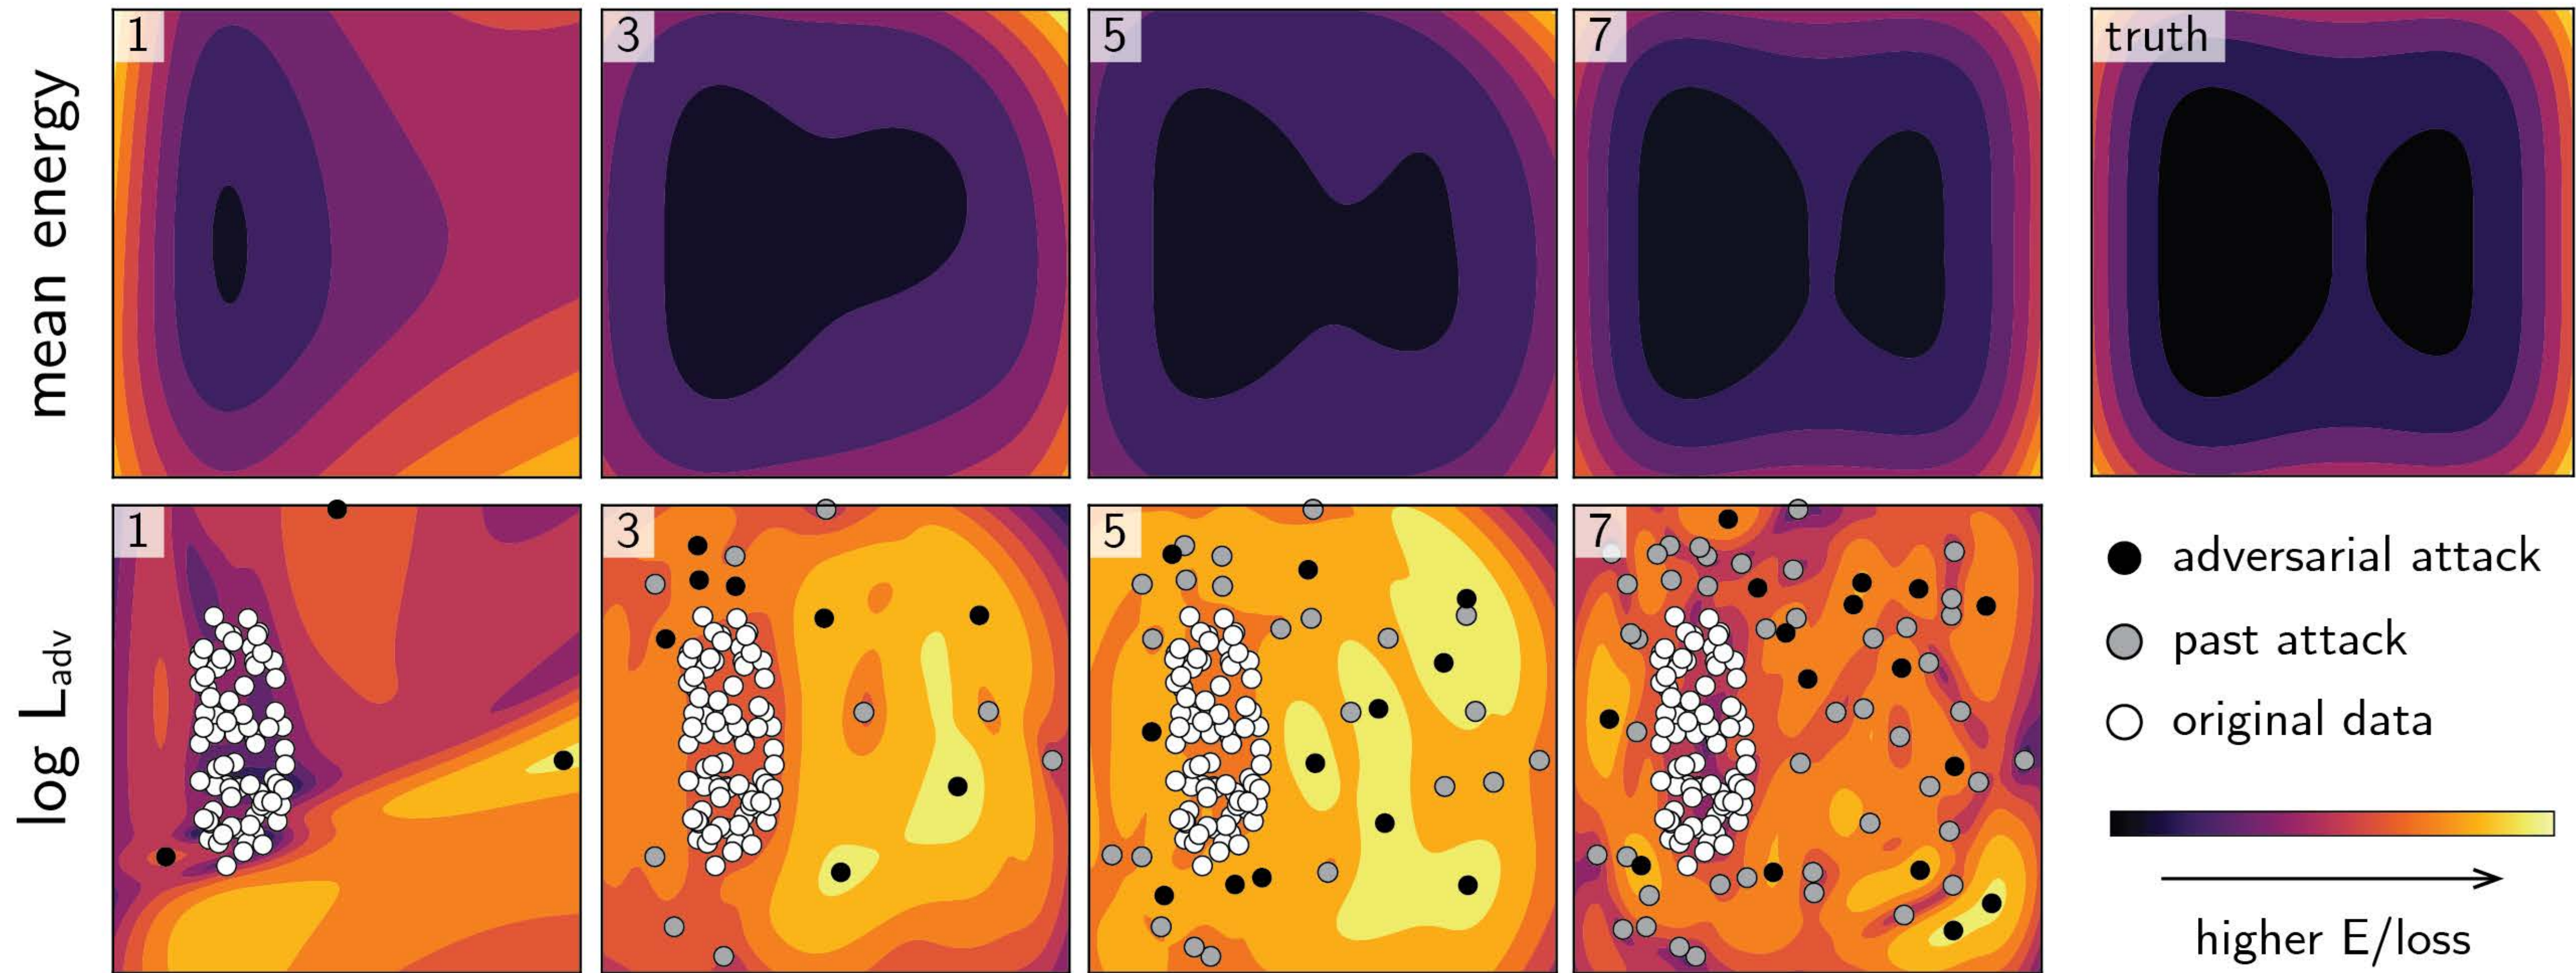
$$p(X_\delta) = \frac{1}{Q} \exp\left(\frac{-\bar{E}(X_\delta)}{kT}\right)$$

The final adversarial objective then becomes

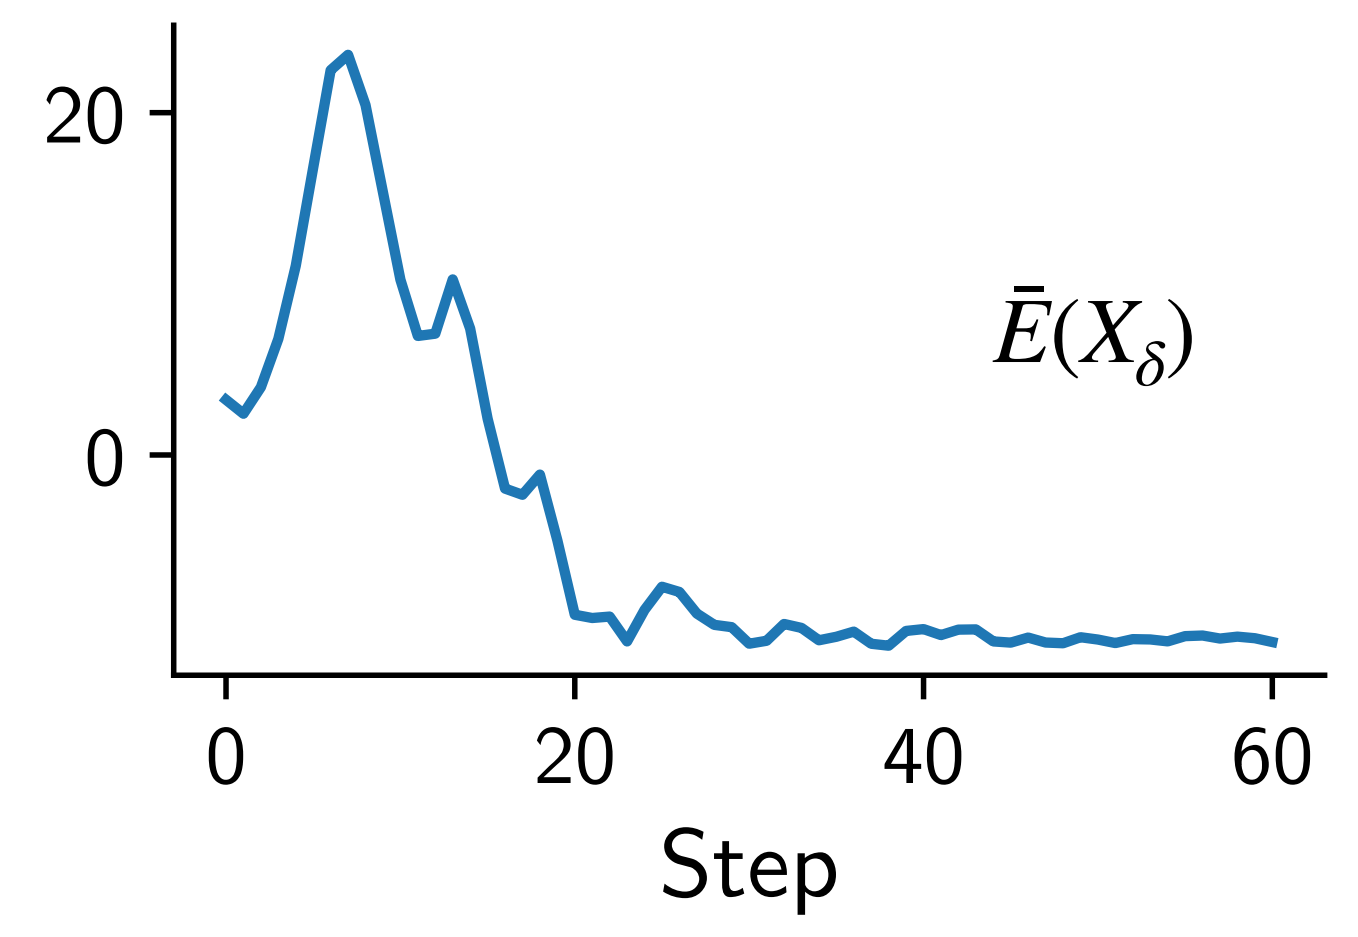




$$\max_{\delta} \mathcal{L}(X, \delta; \theta) = \max_{\delta} p(X_\delta) \cdot \sigma_F^2(X_\delta)$$



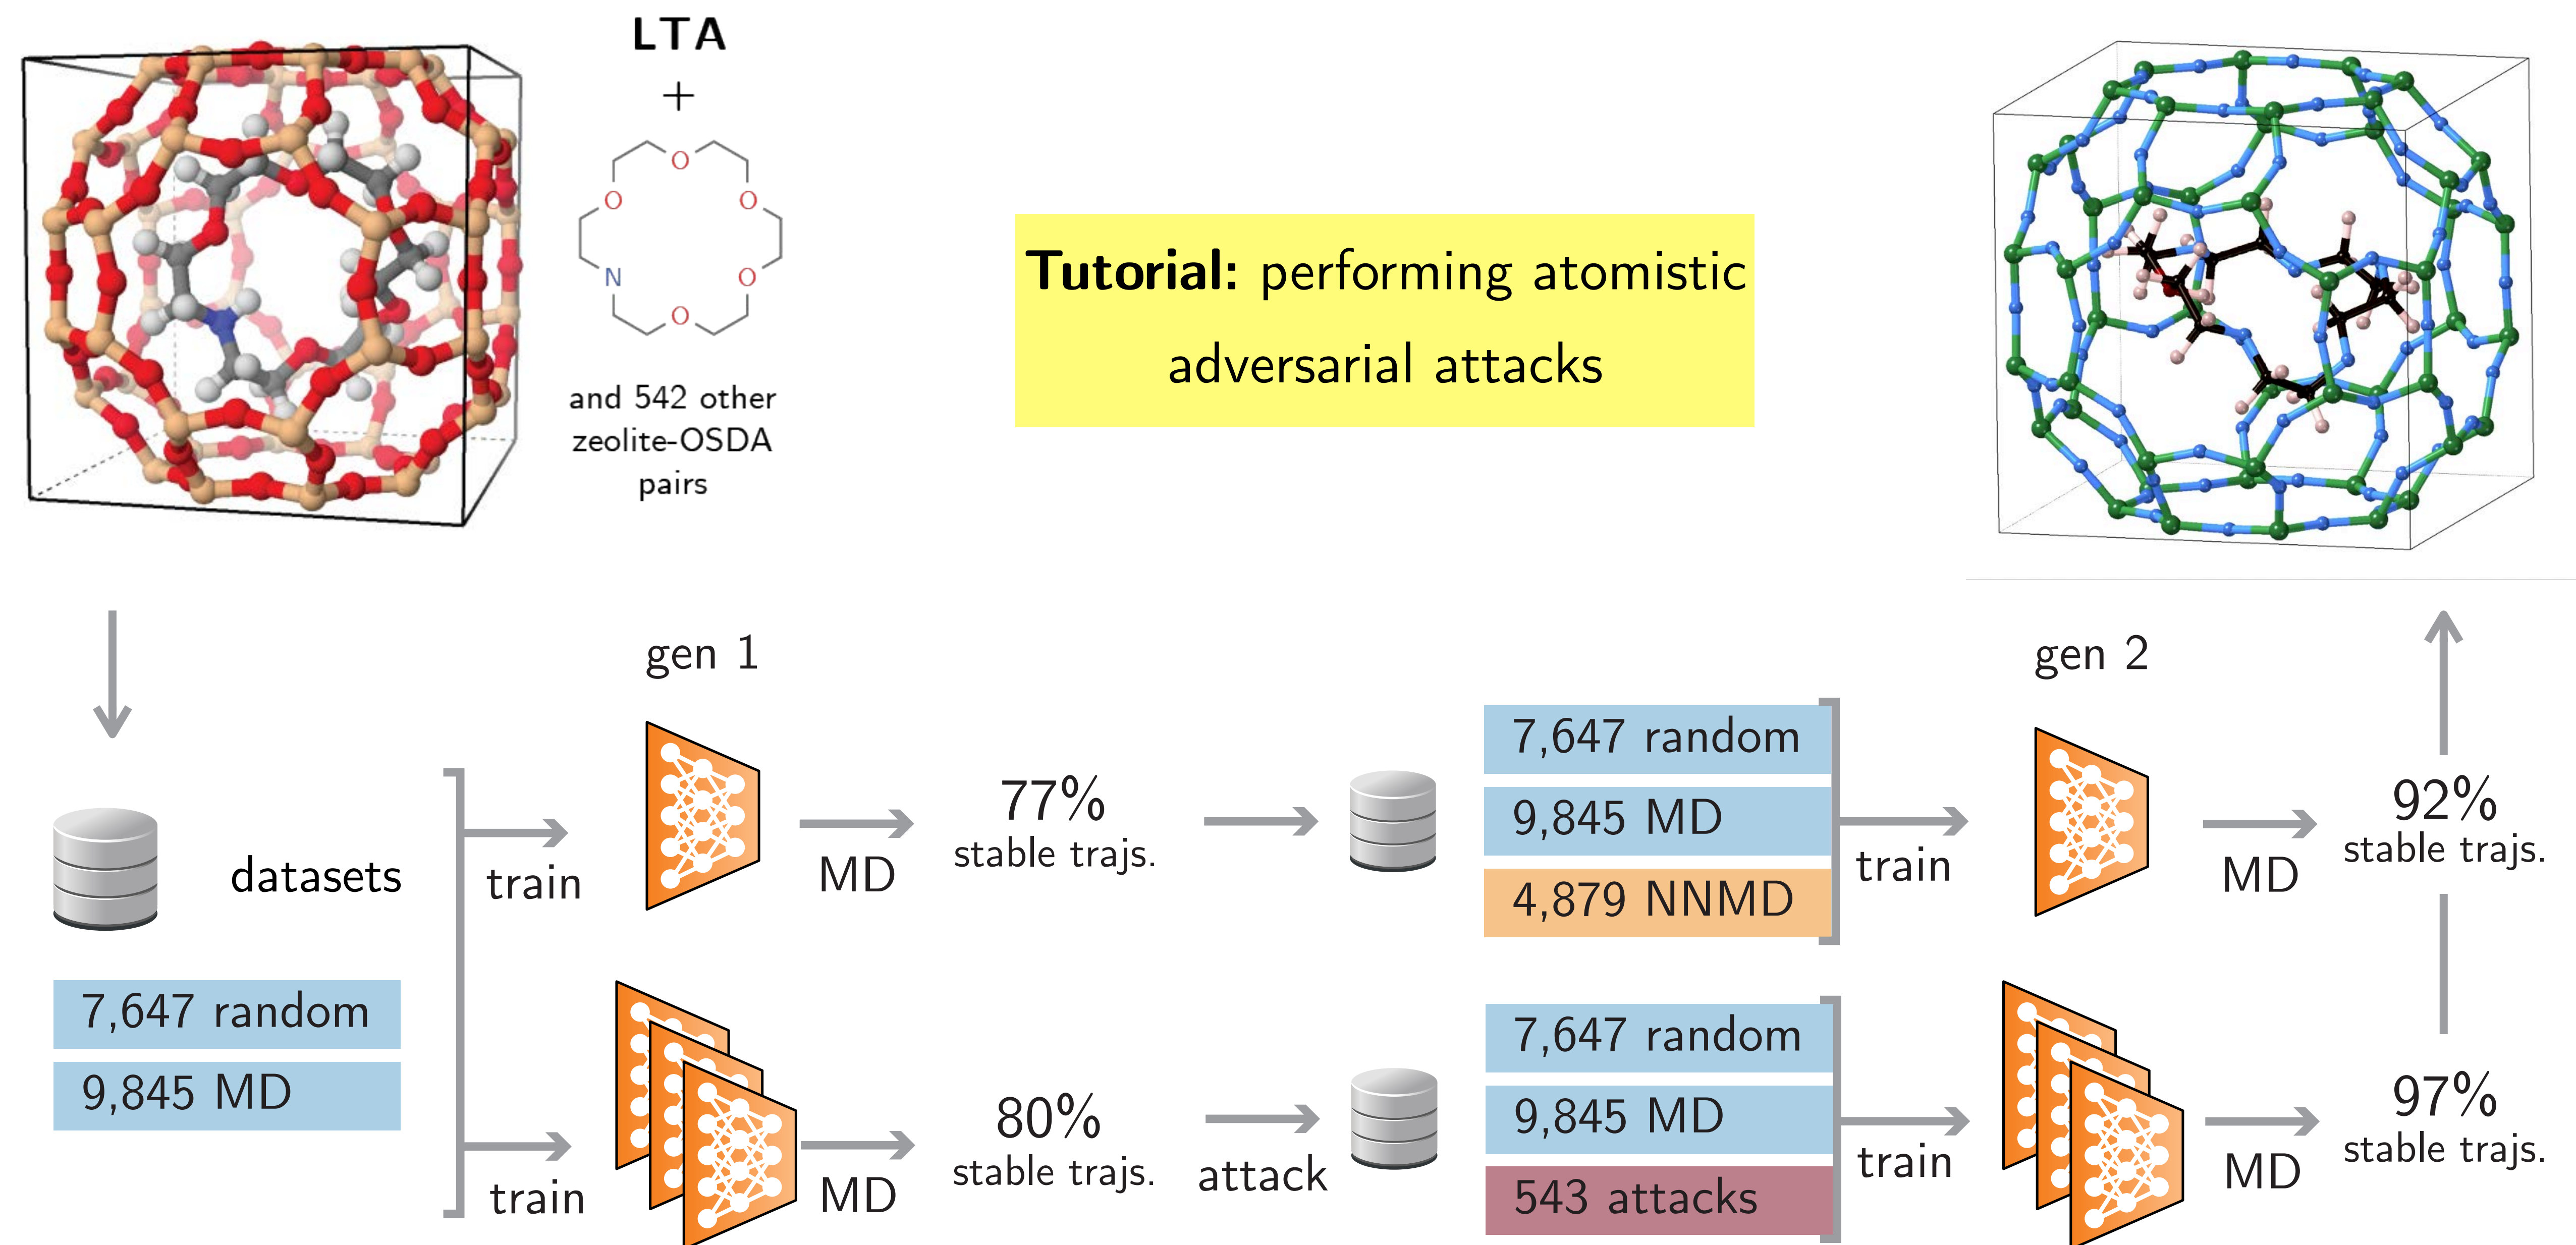
Adversarial attacks for 2D double well



How adversarial attacks look like for molecules?




How efficient is the active learning with this technique?



Summary

What did we learn today?

To summarize what we have learned today



A few resources to learn more

Chemical Reviews **121** (16) (2021): Several reviews on ML for materials

Papers cited in this presentation: In-depth discussion on the advances of NNIPs and much more.

Andrew White's dmol.pub (<https://dmol.pub/>): interactive resources to learn more about ML, deep learning, and their applications to molecules and materials

Michael Nielsen's online book (<http://neuralnetworksanddeeplearning.com/index.html>): several explanations on the math/workings of neural networks

3blue1brown's videos on NNs: excellent visualizations and explanations on NNs
(https://youtube.com/playlist?list=PLZHQObOWTQDNU6R1_67000Dx_ZCJB-3pi)

I. Goodfellow et al. *Deep Learning*. MIT Press (2016): in-depth discussion of deep learning theory (<https://www.deeplearningbook.org/>)

Interatomic Potentials Enabled by Machine Learning

Daniel Schwalbe-Koda

Lawrence Livermore National Laboratory
<https://dskoda.com>, <dskoda@llnl.gov>



Lecture: 4th IKZ-FAIRmat Winter School
Jan. 24, 2023

LLNL-PRES-843937

Prepared by LLNL under Contract DE-AC52-07NA27344. LDRD # 22-ERD-055

Image created with DALL-E 2