Simulation Recipes from the Parrinello Family Cookbook

Michele Ceriotti
cosmo.epfl.ch

Simulating Parrinello-Style

- Know the classics, and use them creatively
- Don't be afraid of complicated s**t
- No project is finished until there's an implementation people can use

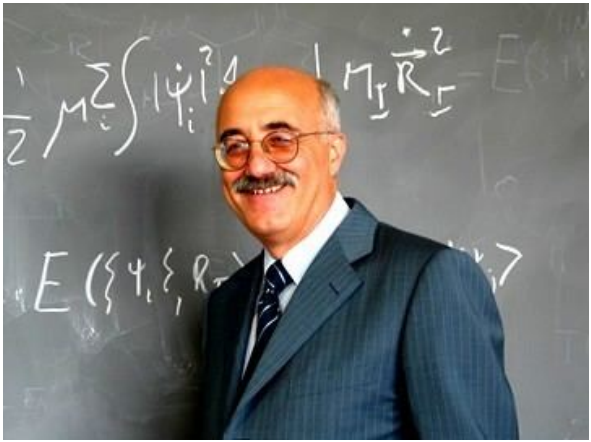

Car & Parrinello, Phys. Rev. Lett. (1985)

Simulating Parrinello-Style

- Know the classics, and use them creatively
- Don't be afraid of complicated s**t
- No project is finished until there's an implementation people can use

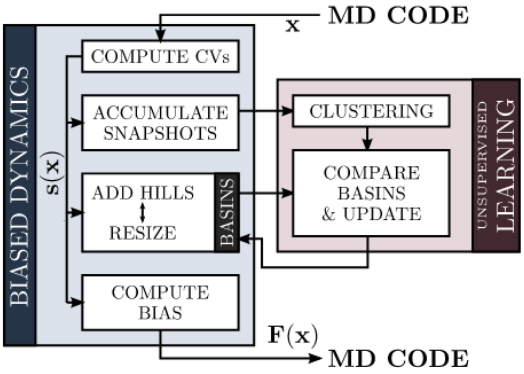

Fig. 2. Flowchart for the reconnaissance metadynamics algorithm.

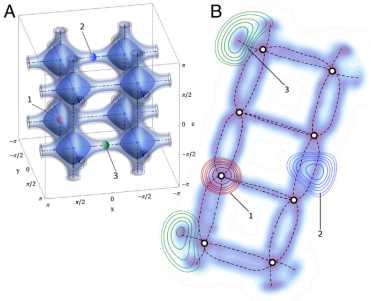

Fig. 1. A complex free-energy surface that is periodic in three directions (*A*) and its sketch-map projection (*B*). *B* shows how one can use functions of the sketch-map coordinates to describe the position in the three-dimensional space. The fields generated for the three marked points are shown. Where sketch-map reproduces the topology (point 1) the field is sharply peaked and is roughly Gaussian shaped. Where sketch-map provides a less good description the field has multiple peaks because there are multiple points where it is reasonable to project (point 3).

Tribello et al. PNAS (2010); Tribello et al. PNAS (2012)

Simulating Parrinello-Style

- Know the classics, and use them creatively
- Don’t be afraid of complicated s**t
- No project is finished until there’s an implementation people can use

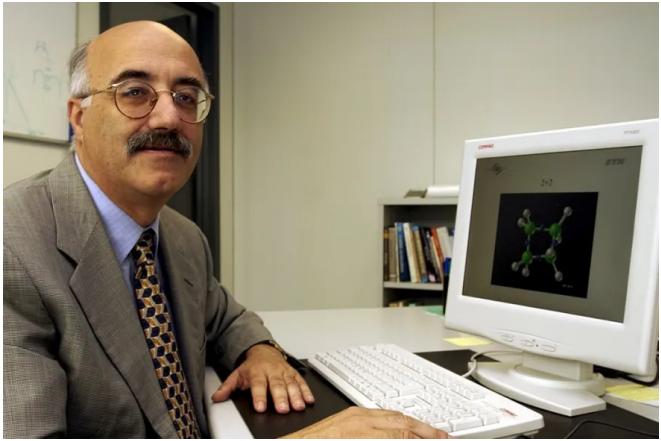

CPMD, CP2K, PLUMED, . . .

A cookbook of atomistic recipes

- Modern atomistic simulations combine QM, stat mech, ML and more
- A comprehensive open-source software stack for advanced atomic-scale simulations
- . . . and a cookbook to use it effectively

molecular dynamics

path integrals

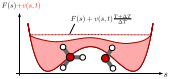

accelerated sampling

wavefunction methods & QMC

multi-scale modeling

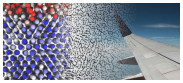

experimental constraints

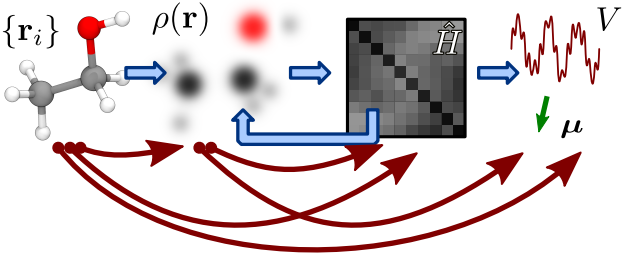

A cookbook of atomistic recipes

- Modern atomistic simulations combine QM, stat mech, ML and more
- A comprehensive open-source software stack for advanced atomic-scale simulations
- . . . and a cookbook to use it effectively

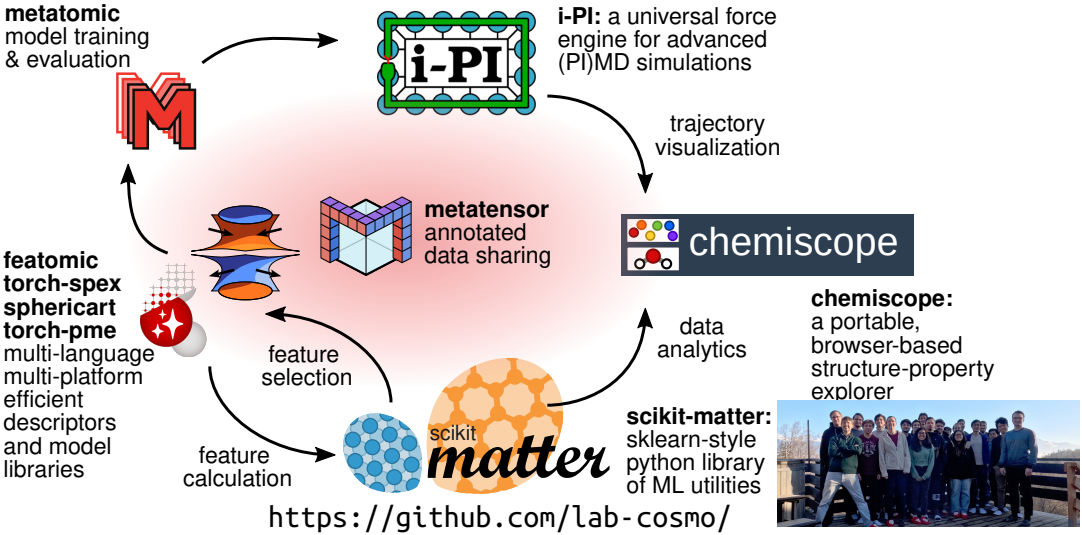

A cookbook of atomistic recipes

- Modern atomistic simulations combine QM, stat mech, ML and more
- A comprehensive open-source software stack for advanced atomic-scale simulations
- . . . and a cookbook to use it effectively

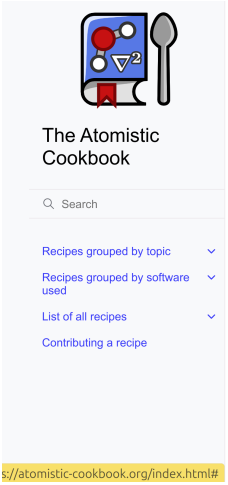

The Atomistic Cookbook

This cookbook contains recipes for atomic-scale modelling for materials and molecules, with a particular focus on machine learning and statistical sampling methods. Most of the examples rely heavily on software developed by the laboratory of computational science and modeling (COSMO, see its github page) but the cookbook is open for recipes using all types of modeling tools and techniques. Rather than focusing on the usage of a specific package, this cookbook provides concrete examples of the solution of modeling problems, often using a combination of several tools.

Downloading and running the recipes

Each recipe can be viewed online as an interactive HTML page, but can also be downloaded as a stand-alone `.py` script of `.ipynb` Jupyter notebook. To simplify setting up an environment that contains all the dependencies needed for each recipe, you can also download an `environment.yml` file that you can use with conda to create a custom environment to run the example.

```
# Pick a name for the environment and replace <environment-name> with it
conda env create --name <environment-name> --file environment.yml

# when you want to use the environment
conda env activate --name <environment-name>
```

Additional data needed for each example is usually either downloaded dynamically, or can be found in a `data` folder for each example, or downloaded as a `data.zip` file at the end of each recipe in the

https://atomistic-cookbook.org/

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture

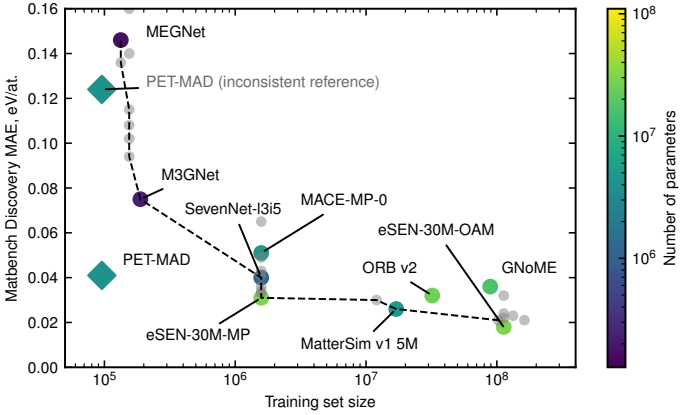

Mazitov

Bigi

Mazitov, Bigi et al., arXiv:2503.14118;
`https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html`

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture

Dataset	PET-MAD	MACE-MP-0-L	MatterSim-5M	Orb-v2	SevenNet
MAD	**17.6**|**65.1**	81.6|181.5	47.3|133.7	52.9|96.2	82.1|173.5
MPtrj	22.3|77.9	15.1|50.8	21.3|61.4	**5.6**|**21.9**	9.8|25.5
Matbench	**31.3**| ---	58.5|---	38.2 |---	37.9|---	47.5|---
Alexandria	49.0|66.8	65.4|79.5	21.2|39.9	**13.2**|**10.5**	47.6|70.3
OC2020	**18.3**|114.5	82.4|169.6	31.5|119.2	19.8|**99.3**	45.7|162.7
SPICE	**3.7**|**59.5**	10.6|166.8	21.3|145.6	59.0|140.8	11.3|139.1
MD22	**1.9**|**65.6**	9.4|182.9	28.6|160.4	174.3|220.7	11.1|146.2

Mazitov, Bigi et al., arXiv:2503.14118;
`https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html`

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture

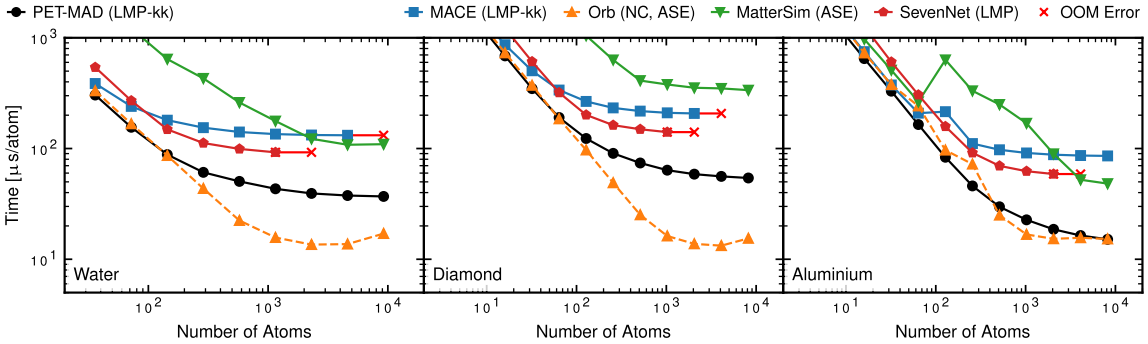

Mazitov, Bigi et al., arXiv:2503.14118;
`https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html`

atomistic-cookbook.org/examples/pet-mad-uq/pet-mad-uq.html

Go to the end to download the full example code.

Uncertainty Quantification with PET-MAD

Authors:

Johannes Spies @johannes-spies

This recipe demonstrates three ways of computing errors on the outputs of ML potential-driven simulations, using as an example the PET-MAD model and its built-in uncertainty quantification (UQ) capabilities.

In particular, it demonstrates:

1. Estimating uncertainties for single-point calculations on a full validation dataset.
2. Computing energies in simple functions of energy predictions, namely the value of vacancy formation energies
3. Propagating errors from energy predictions to thermodynamic averages computed over a constant-temperature MD simulation.

For more information on PET-MAD, have a look at Mazitov et al., 2025. The LLPR uncertainties are introduced in Bigi et al., 2024. For more information on dataset calibration and error propagation, see Imabalzano et al., 2021.

Optional: Adding UQ to a Model

ON THIS PAGE

The design space of uncertainty quantification

- Large design space for uncertainty quantification on top of a given architecture:
  1. Error estimate (ensemble vs mean-variance dual-output)
  2. Diversification of the ensemble (subsampling, weight pruning, NLL target)
  3. Calibration (essential for quantitative estimates)
  4. Uncertainty propagation

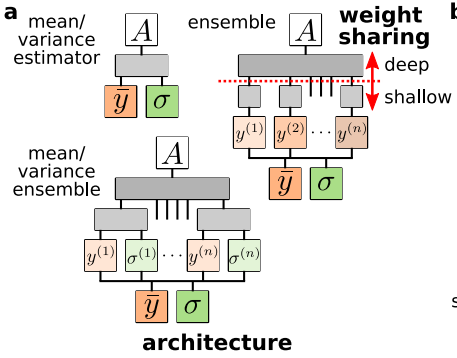

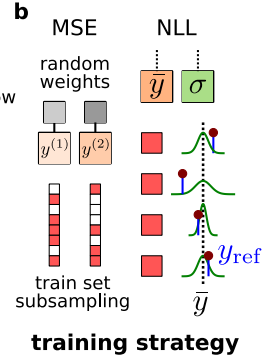

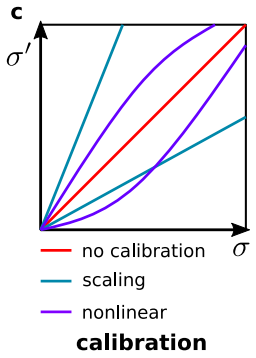

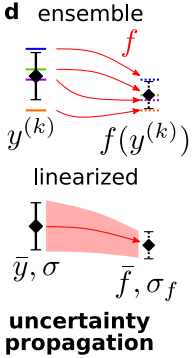

Uncertainty estimation on a shoestring

- "Shallow" ensemble architecture to reduce overhead to (nearly) zero
- Propagate the ensemble directly to obtain errors on the end properties

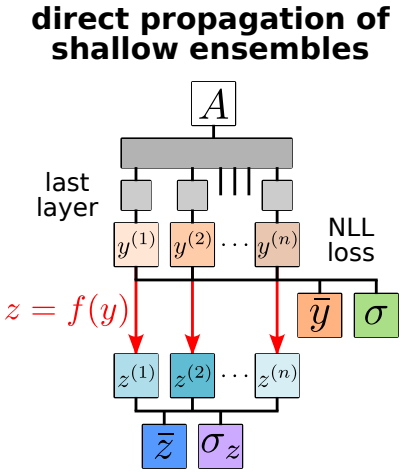

Musil et al., JCTC (2019); Kellner & MC, MLST (2024)

Last-layer ensembles are accurate

- A last-layer ensemble model provides good, calibrated and informative UQ over a variety of atomistic benchmarks

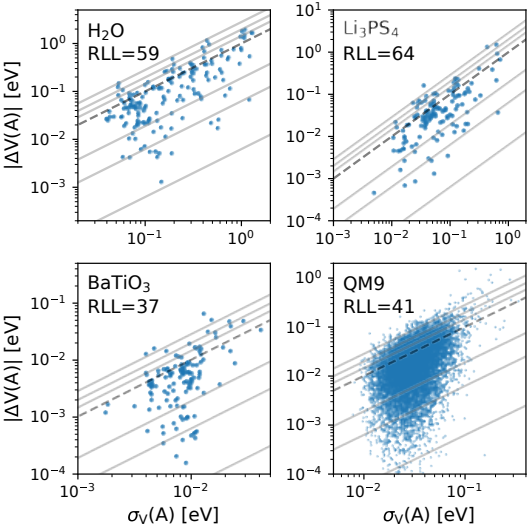

Kellner & MC, MLST (2024)

Ensembles are easy to propagate uncertainty with

- Obtaining uncertainties for complex simulation protocols is as simple as applying them to separate ensemble members

$$z^{(k)}(A) = f\left[y^{(k)}(A)\right] \qquad \bar{z} = \frac{1}{n}\sum_k z^{(k)},\ \sigma_z^2 = \frac{1}{n-1}\sum_k \left(z^{(k)} - \bar{z}\right)^2$$

	mean / eV	var / eV	std / eV
Before creating vacancy	-127.312576	1.347783	1.160941
Right after creating vacancy	-122.585747	1.116814	1.056794
Energy of optimized vacancy	-122.649384	1.129488	1.062774
Vacancy formation energy	0.684676	0.011522	0.107341

Ensembles are easy to propagate uncertainty with

- Obtaining uncertainties for complex simulation protocols is as simple as applying them to separate ensemble members

$$z^{(k)}(A) = f\left[y^{(k)}(A)\right] \qquad \bar{z} = \frac{1}{n}\sum_k z^{(k)},\ \sigma_z^2 = \frac{1}{n-1}\sum_k \left(z^{(k)} - \bar{z}\right)^2$$

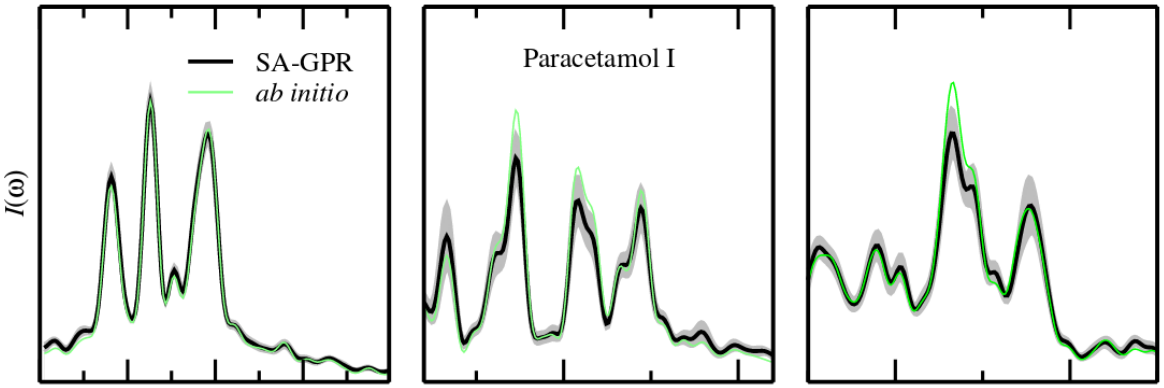

N. Raimbault, A. Grisafi, MC, M. Rossi, New J. Phys. (2019)

Ensembles are easy to propagate uncertainty with

- Obtaining uncertainties for complex simulation protocols is as simple as applying them to separate ensemble members

$$z^{(k)}(A) = f\left[y^{(k)}(A)\right] \qquad \bar{z} = \frac{1}{n}\sum_k z^{(k)},\ \sigma_z^2 = \frac{1}{n-1}\sum_k \left(z^{(k)} - \bar{z}\right)^2$$

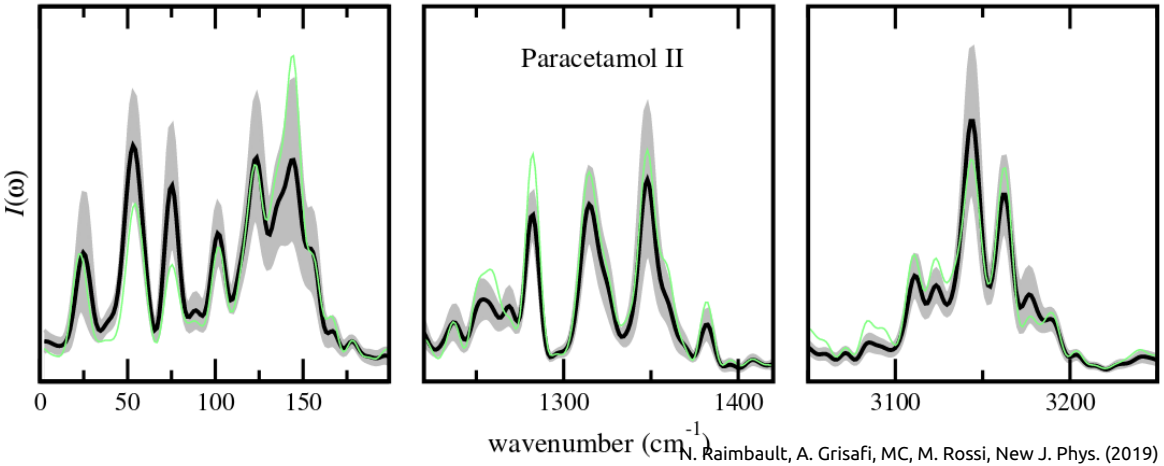

N. Raimbault, A. Grisafi, MC, M. Rossi, New J. Phys. (2019)

UQ for free-energy sampling

- Errors in ML-based thermodynamic averages combine effects on the observable σ_a^2 and those from sampling σ_{aV}^2
- A committee of predictions can be obtained from a single trajectory!

$$\langle a \rangle_{V^{(i)}} = \left\langle a\, e^{\beta(\bar{V}-V^{(i)})} \right\rangle_{\bar{V}}$$

- Statistically stable estimates with a Cumulant Expansion Approximation

$$\langle a \rangle_{V^{(i)},\mathrm{CEA}} \approx \langle a \rangle_{\bar{V}} - \beta \left[\left\langle a \left(V^{(i)} - \bar{V} \right) \right\rangle_{\bar{V}} - \langle a \rangle_{\bar{V}} \left\langle \left(V^{(i)} - \bar{V} \right) \right\rangle_{\bar{V}} \right]$$

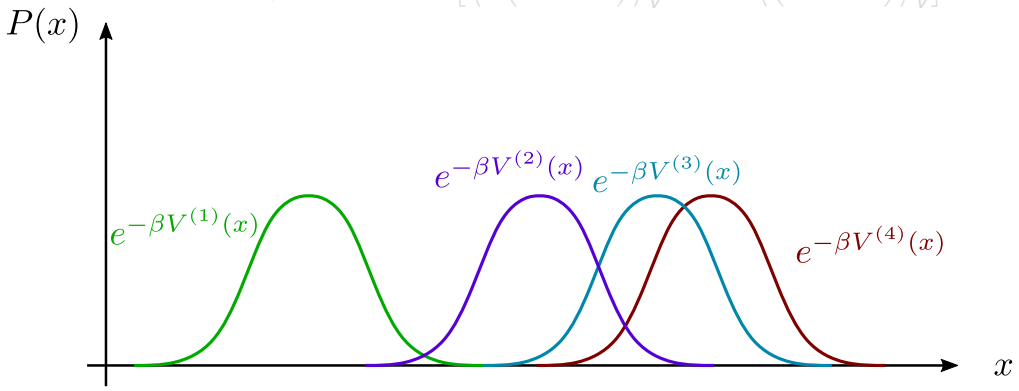

Imbalzano et al. JCP (2021)

UQ for free-energy sampling

- Errors in ML-based thermodynamic averages combine effects on the observable σ_a^2 and those from sampling σ_{aV}^2
- A committee of predictions can be obtained from a single trajectory!

$$\langle a \rangle_{V^{(i)}} = \left\langle a\, e^{\beta(\bar{V}-V^{(i)})} \right\rangle_{\bar{V}}$$

- Statistically stable estimates with a Cumulant Expansion Approximation

$$\langle a \rangle_{V^{(i)},\mathrm{CEA}} \approx \langle a \rangle_{\bar{V}} - \beta \left[\left\langle a \left(V^{(i)} - \bar{V} \right) \right\rangle_{\bar{V}} - \langle a \rangle_{\bar{V}} \left\langle \left(V^{(i)} - \bar{V} \right) \right\rangle_{\bar{V}} \right]$$

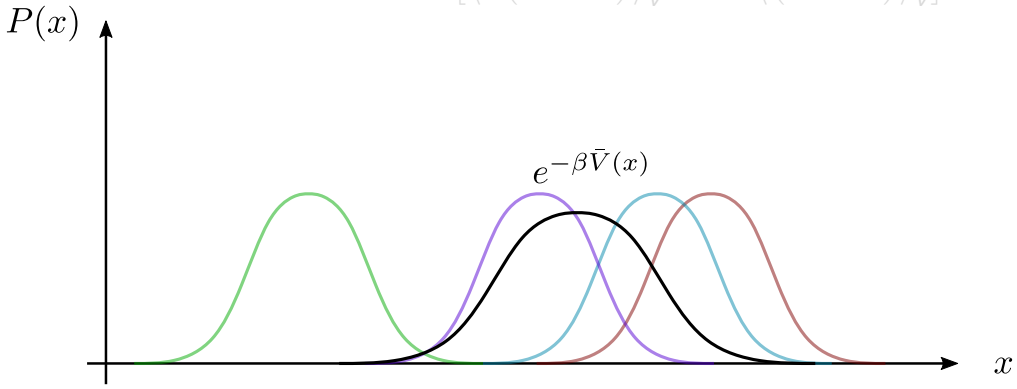

Imbalzano et al. JCP (2021)

UQ for free-energy sampling

- Errors in ML-based thermodynamic averages combine effects on the observable σ_a^2 and those from sampling σ_{aV}^2
- A committee of predictions can be obtained from a single trajectory!

$$\langle a \rangle_{V^{(i)}} = \left\langle a\, e^{\beta(\bar{V}-V^{(i)})} \right\rangle_{\bar{V}}$$

- Statistically stable estimates with a Cumulant Expansion Approximation

$$\langle a \rangle_{V^{(i)},\mathrm{CEA}} \approx \langle a \rangle_{\bar{V}} - \beta \left[\left\langle a \left(V^{(i)} - \bar{V} \right) \right\rangle_{\bar{V}} - \langle a \rangle_{\bar{V}} \left\langle \left(V^{(i)} - \bar{V} \right) \right\rangle_{\bar{V}} \right]$$

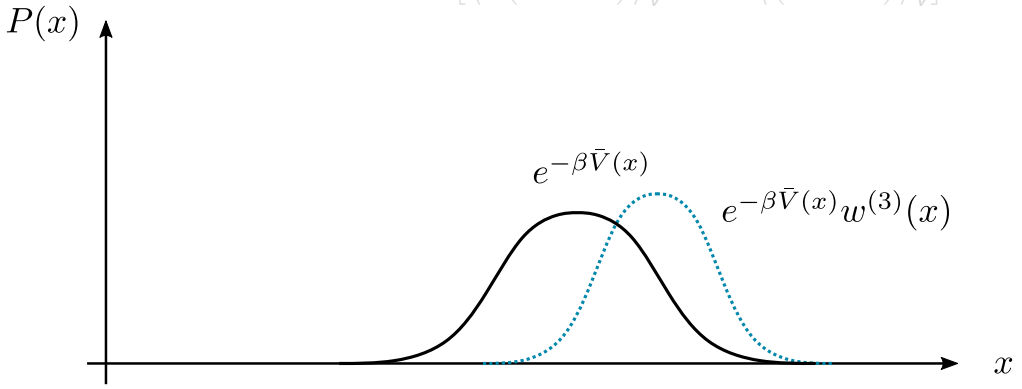

Imbalzano et al. JCP (2021)

UQ for free-energy sampling

- Errors in ML-based thermodynamic averages combine effects on the observable σ_a^2 and those from sampling σ_{aV}^2
- A committee of predictions can be obtained from a single trajectory!

$$\langle a \rangle_{V^{(i)}} = \left\langle a\, e^{\beta(\bar{V}-V^{(i)})} \right\rangle_{\bar{V}}$$

- Statistically stable estimates with a Cumulant Expansion Approximation

$$\langle a \rangle_{V^{(i)},\mathrm{CEA}} \approx \langle a \rangle_{\bar{V}} - \beta \left[\left\langle a \left(V^{(i)} - \bar{V} \right) \right\rangle_{\bar{V}} - \langle a \rangle_{\bar{V}} \left\langle \left(V^{(i)} - \bar{V} \right) \right\rangle_{\bar{V}} \right]$$

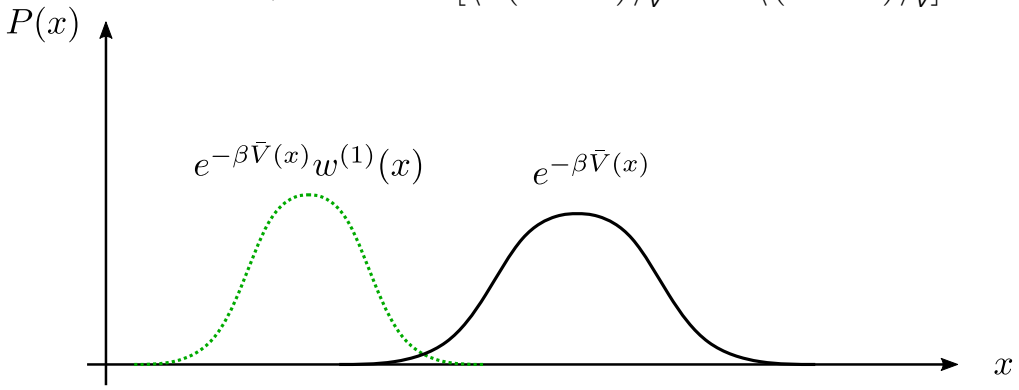

MC et al. PRSA (2012)

UQ propagation with PET-MAD

- Verify accuracy of a universal model on the end property by comparison with a bespoke model or a fine tuned one)
- Predictive model errors with shallow ensembles

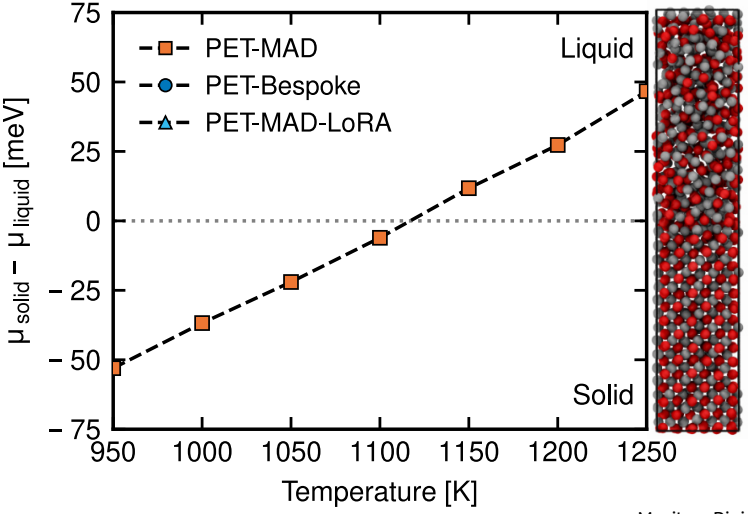

Mazitov, Bigi et al., arXiv:2503.14118;

UQ propagation with PET-MAD

- Verify accuracy of a universal model on the end property by comparison with a bespoke model or a fine tuned one)
- Predictive model errors with shallow ensembles

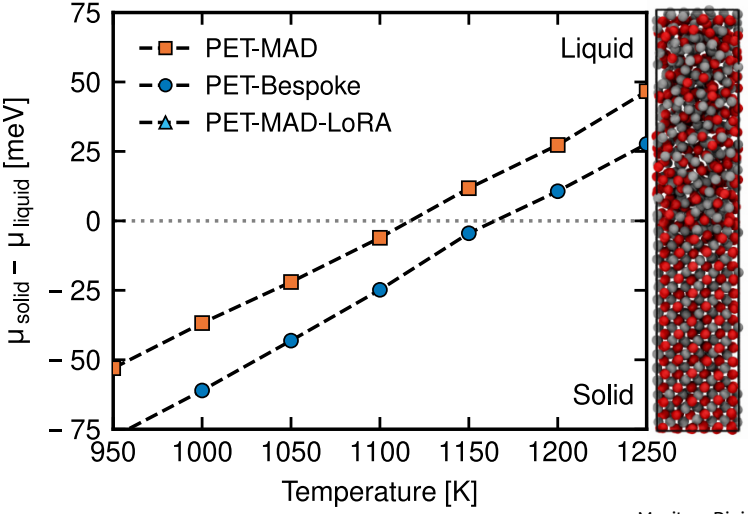

Mazitov, Bigi et al., arXiv:2503.14118;

UQ propagation with PET-MAD

- Verify accuracy of a universal model on the end property by comparison with a bespoke model or a fine tuned one)
- Predictive model errors with shallow ensembles

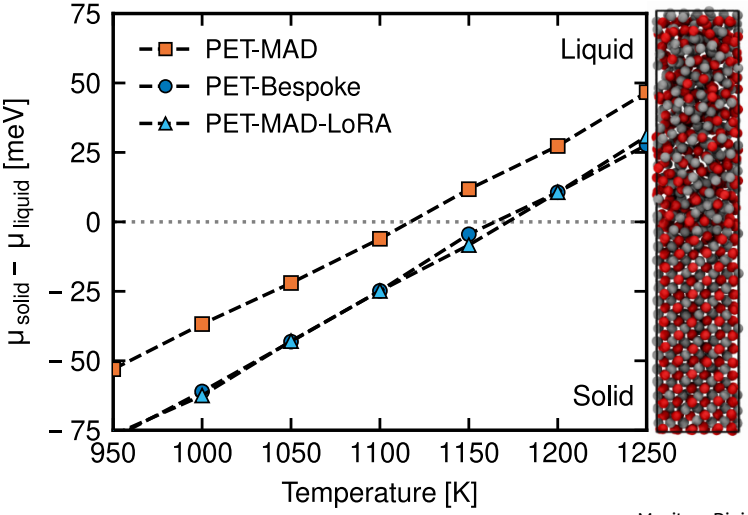

Mazitov, Bigi et al., arXiv:2503.14118;

UQ propagation with PET-MAD

- Verify accuracy of a universal model on the end property by comparison with a bespoke model or a fine tuned one)
- Predictive model errors with shallow ensembles

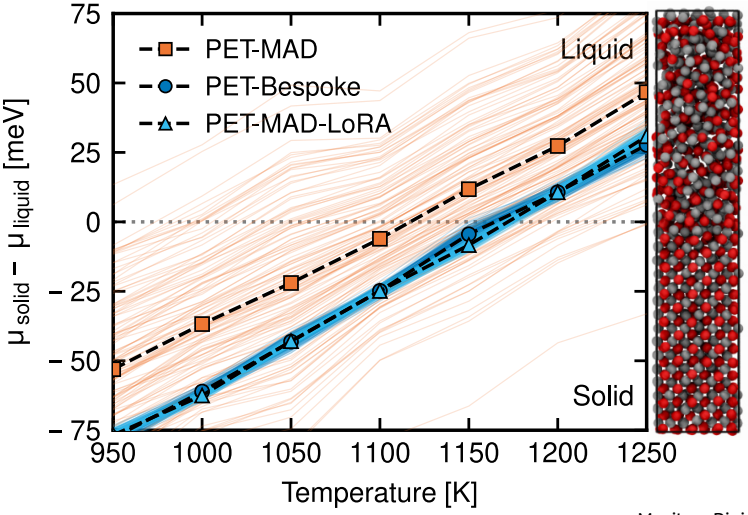

Mazitov, Bigi et al., arXiv:2503.14118;

atomistic-cookbook.org/examples/pet-mad-nc/pet-mad-nc.html

Go to the end to download the full example code.

MD using direct-force predictions with PET-MAD

Authors:

Michele Ceriotti @ceriottm, Filippo Bigi @frostedoyster

Evaluating forces as a direct output of a ML model accelerates their evaluation, by up to a factor of 3 in comparison to the traditional approach that evaluates them as derivatives of the interatomic potential. Unfortunately, as discussed e.g. in this paper, doing so means that forces are not conservative, leading to instabilities and artefacts in many modeling tasks, such as constant-energy molecular dynamics. Here we demonstrate the issues associated with direct force predictions, and ways to mitigate them, using the generally-applicable PET-MAD potential. See also this recipe for examples of using PET-MAD for basic tasks such as geometry optimization and conservative MD, and this for an example of how to use direct forces to accelerate training.

```
# sphinx_gallery_thumbnail_number = 2
```

If you don't want to use the conda environment for this recipe, you can get all dependencies installing the PET-MAD package:

```
pip install pet-mad
```

ON THIS PAGE

One (equivariant) framework to rule them all

- Deep connections between most atomic-scale ML frameworks due to symmetry priors
- All equivariant frameworks have a structure constrained by the $O(3)$ group

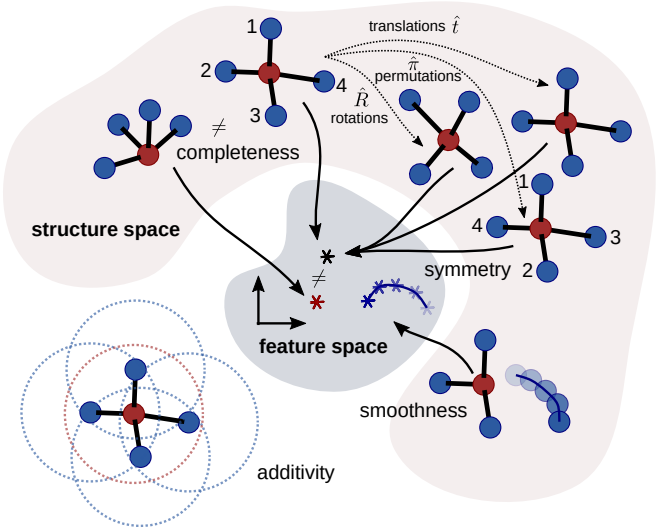


Musil


Nigam

Musil et al., Chem. Rev. (2021)

One (equivariant) framework to rule them all

- Deep connections between most atomic-scale ML frameworks due to symmetry priors
- All equivariant frameworks have a structure constrained by the $O(3)$ group

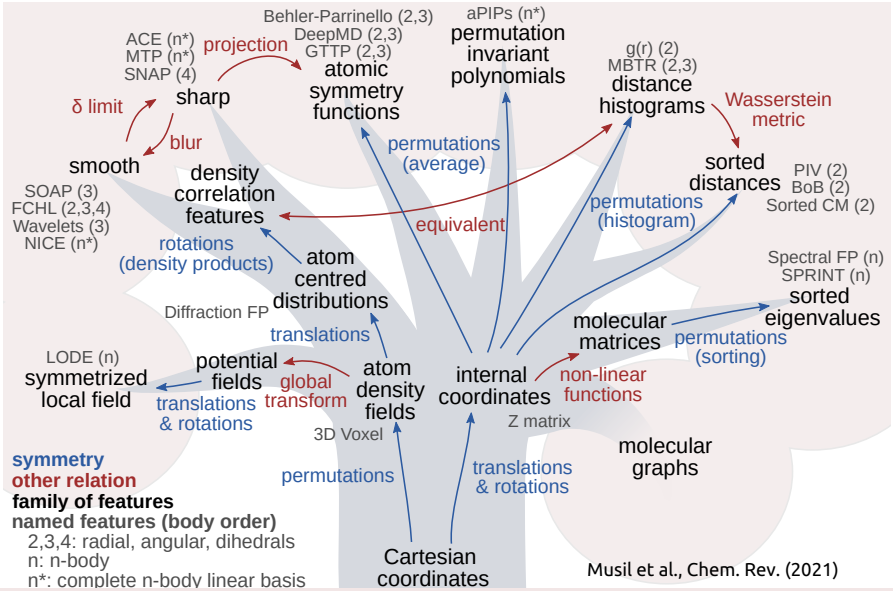

Musil

Nigam

One (equivariant) framework to rule them all

- Deep connections between most atomic-scale ML frameworks due to symmetry priors
- All equivariant frameworks have a structure constrained by the $O(3)$ group

$$\xi_n^{(\lambda\mu)}(A_i) \equiv \langle n|A_i;\lambda\mu\rangle \equiv \sum_{j\in A_i} R_{n\lambda}(r_{ij})\, Y_\lambda^\mu(\hat{\mathbf{r}}_{ij})$$

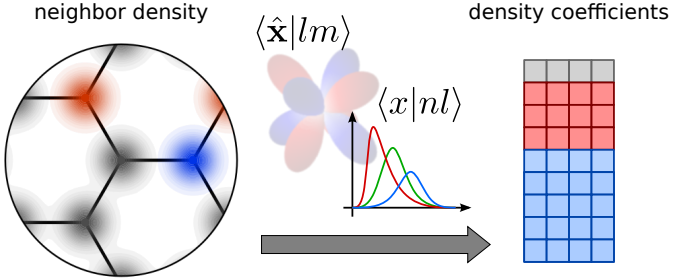


Musil


Nigam

Nigam et al., J. Chem. Phys. (2022)

One (equivariant) framework to rule them all

- Deep connections between most atomic-scale ML frameworks due to symmetry priors
- All equivariant frameworks have a structure constrained by the $O(3)$ group

$$\langle \dots ; n_\nu l_\nu k_\nu ; nlk | \overline{\rho_i^{\otimes(\nu+1)}} ; \lambda\mu \rangle = \sum_{qm} \langle n | \overline{\rho_i^{\otimes 1}} ; lm \rangle \, \langle \dots ; n_\nu l_\nu k_\nu | \overline{\rho_i^{\otimes \nu}} ; kq \rangle \, \langle lm ; kq | \lambda\mu \rangle$$

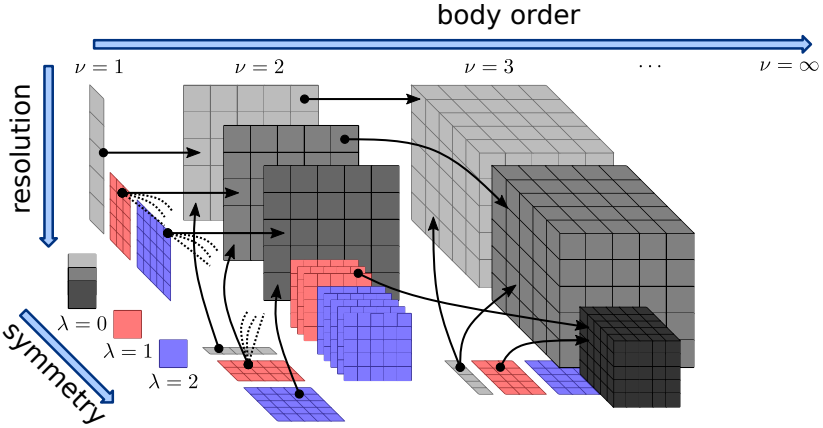

Nigam et al., J. Chem. Phys. (2022)

Musil

Nigam

The bitter lesson

- Deep learning with rotational augmentation is simpler, faster, and often more accurate

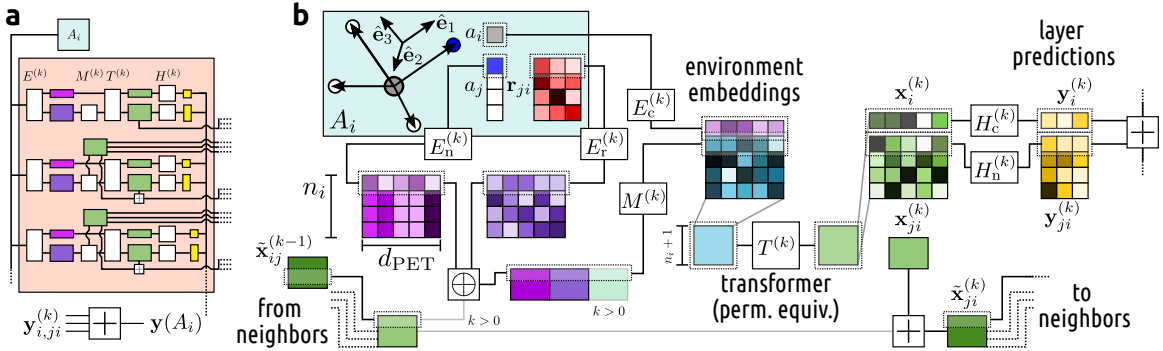

Pozdnyakov, MC, NeurIPS (2023)

The bitter lesson

- Deep learning with rotational augmentation is simpler, faster, and often more accurate

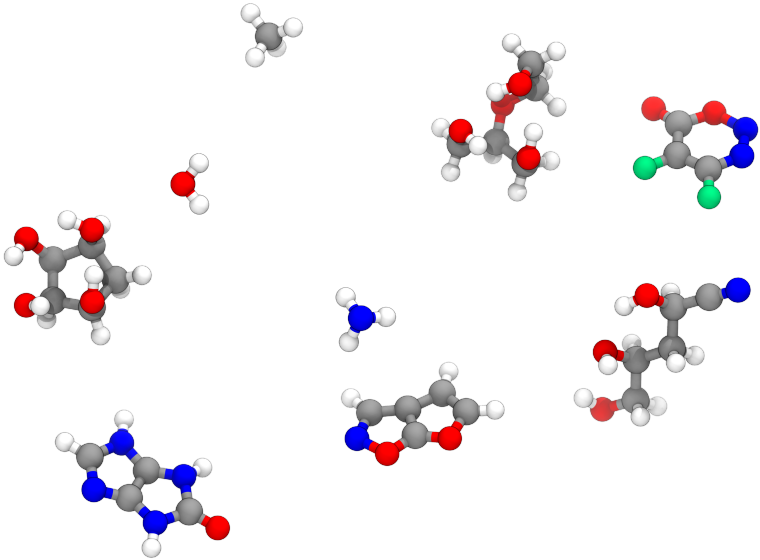

Pozdnyakov, MC, NeurIPS (2023)

The bitter lesson

- Deep learning with rotational augmentation is simpler, faster, and often more accurate

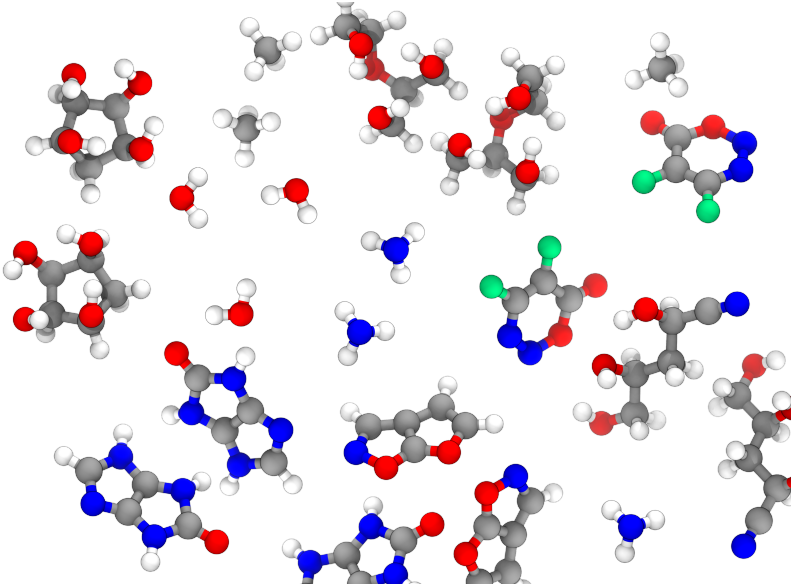

Pozdnyakov, MC, NeurIPS (2023)

The bitter lesson

- Deep learning with rotational augmentation is simpler, faster, and often more accurate

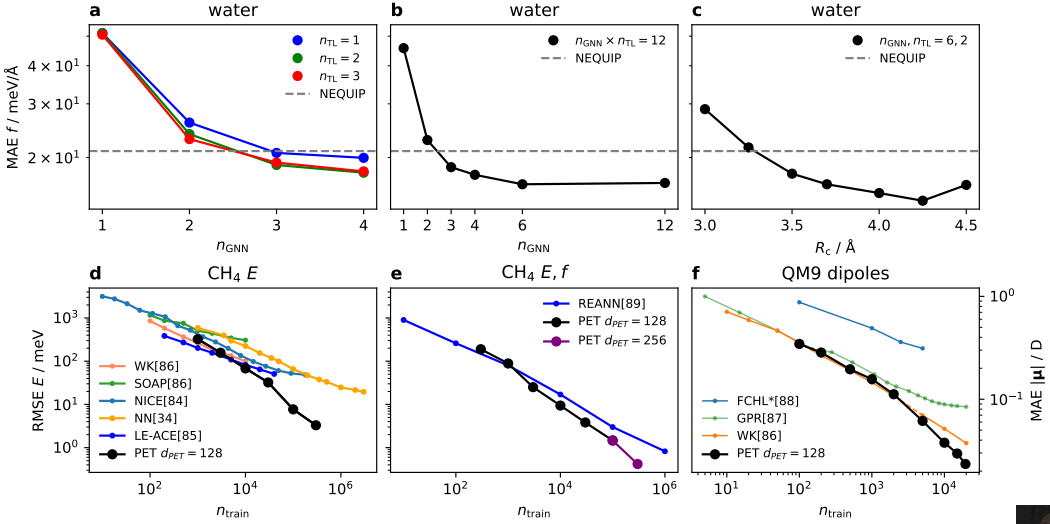

Pozdnyakov, MC, NeurIPS (2023)

The bitter lesson

- Deep learning with rotational augmentation is simpler, faster, and often more accurate

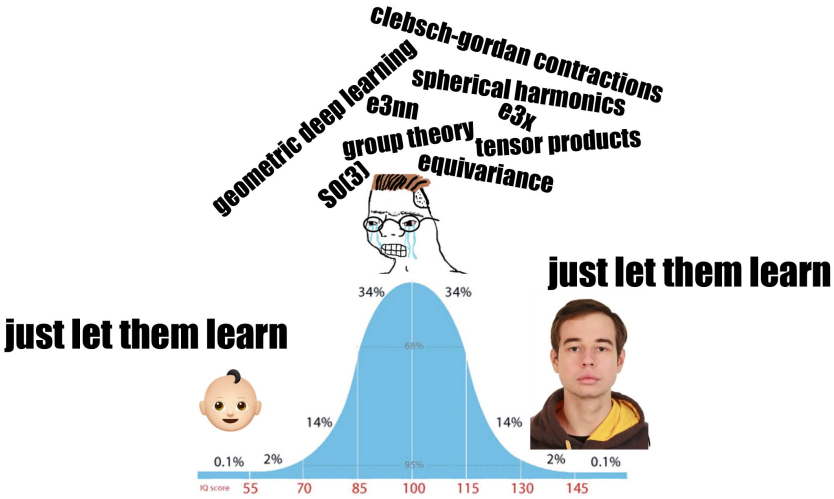

credits: M. Langer

Handle with some care

- Symmetry breaking leads to artefacts

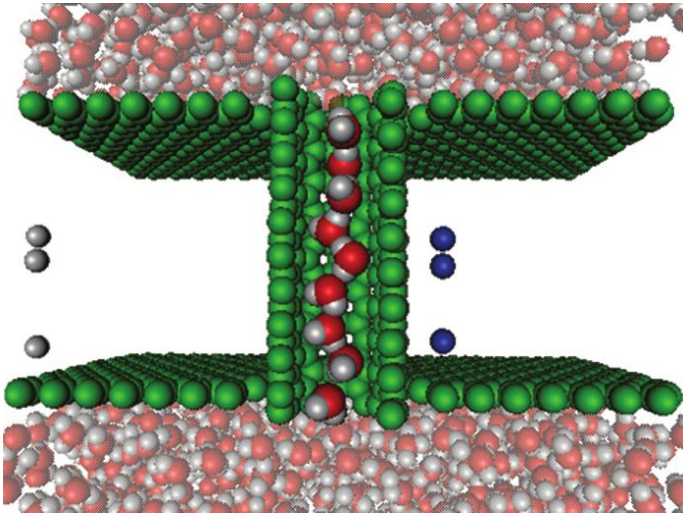

Gong et al., Nature Nano. (2007)

Handle with some care

- Symmetry breaking leads to artefacts that are small and easy to control

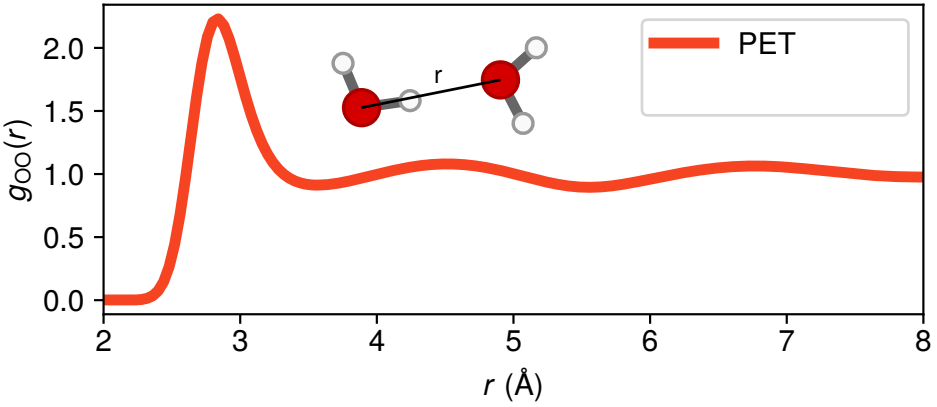

Langer

Pozdnyakov

Langer, Pozdnyakov, MC, MLST (2024)

Handle with some care

- Symmetry breaking leads to artefacts that are small and easy to control

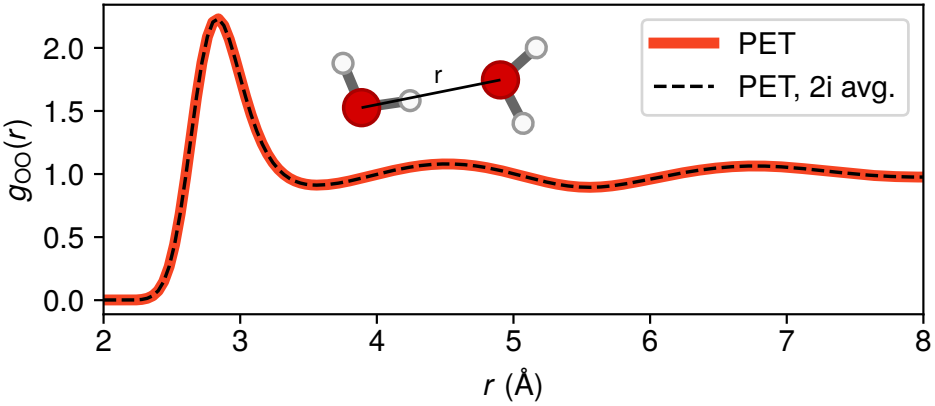

Langer

Pozdnyakov

Langer, Pozdnyakov, MC, MLST (2024)

Handle with some care

- Symmetry breaking leads to artefacts that are small and easy to control

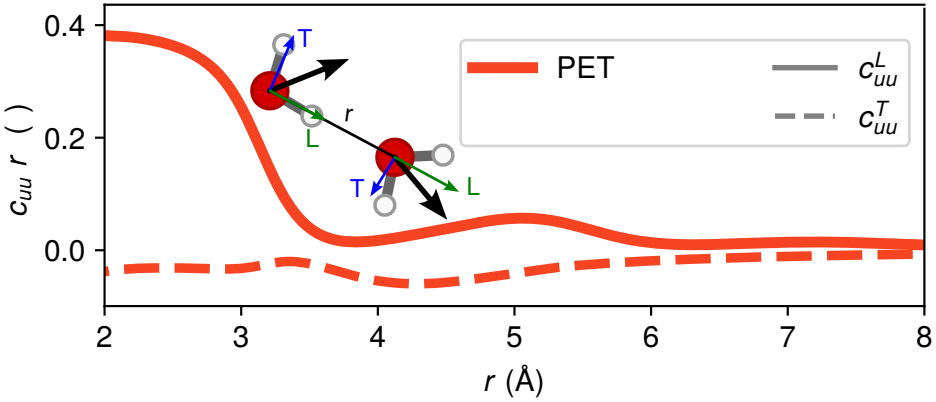

Langer

Pozdnyakov

Langer, Pozdnyakov, MC, MLST (2024)

Handle with some care

- Symmetry breaking leads to artefacts that are small and easy to control

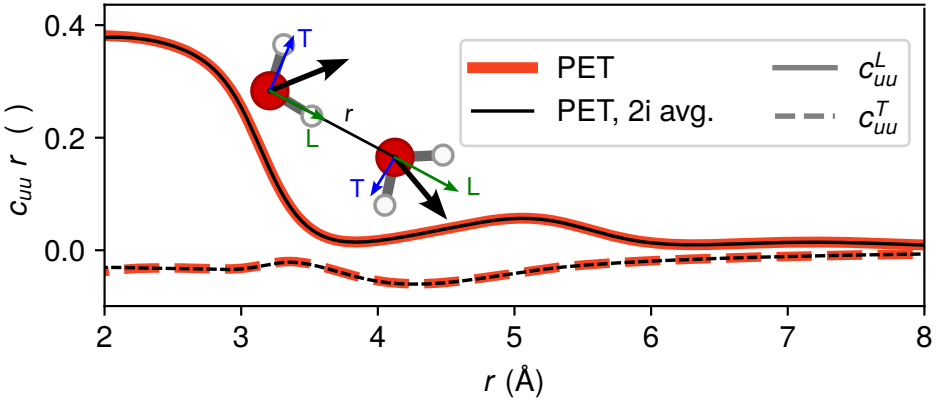

Langer

Pozdnyakov

Langer, Pozdnyakov, MC, MLST (2024)

The dark side of the forces

- Many application only need interatomic *forces*, $\mathbf{f}_j = -\partial V\left(\mathbf{r}_1, \dots \mathbf{r}_N\right) / \partial \mathbf{r}_j$, but learn them as the derivative of a potential energy V
- Data-minded approach: learning $\mathbf{f}_j$ directly is faster
- Bad idea. No energy conservation, unstable sampling, dynamics and geometry optimization

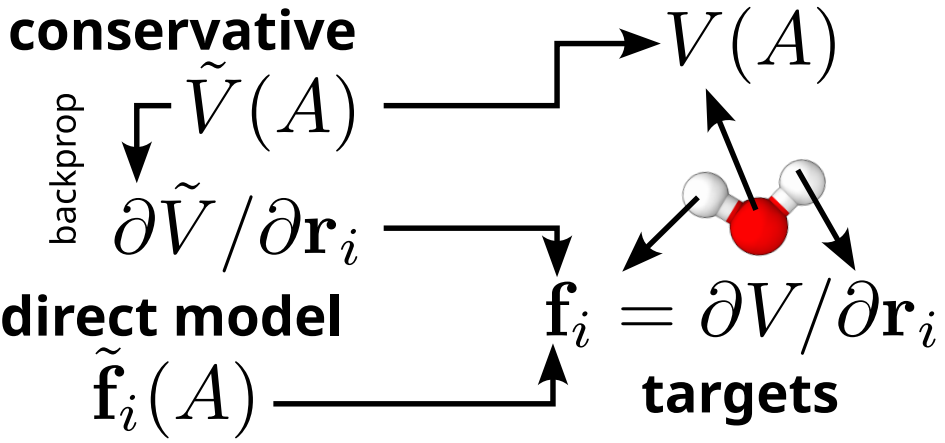


Bigi


Langer

Bigi, Langer, MC, oral @ ICML (2025)

The dark side of the forces

- Many application only need interatomic *forces*, $\mathbf{f}_j = -\partial V\left(\mathbf{r}_1, \ldots \mathbf{r}_N\right) / \partial \mathbf{r}_j$, but learn them as the derivative of a potential energy V
- Data-minded approach: learning $\mathbf{f}_j$ directly is faster
- Bad idea. No energy conservation, unstable sampling, dynamics and geometry optimization

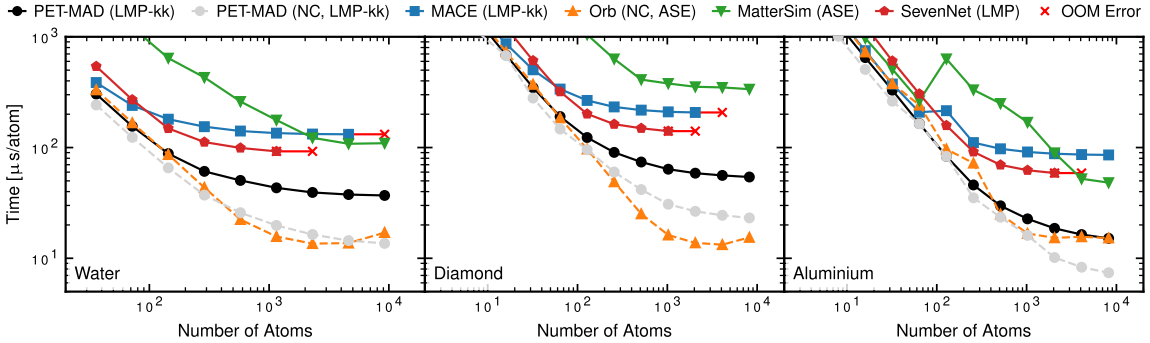

The dark side of the forces

- Many application only need interatomic *forces*, $\mathbf{f}_j = -\partial V\left(\mathbf{r}_1, \ldots \mathbf{r}_N\right) / \partial \mathbf{r}_j$, but learn them as the derivative of a potential energy V
- Data-minded approach: learning $\mathbf{f}_j$ directly is faster
- Bad idea. No energy conservation, unstable sampling, dynamics and geometry optimization

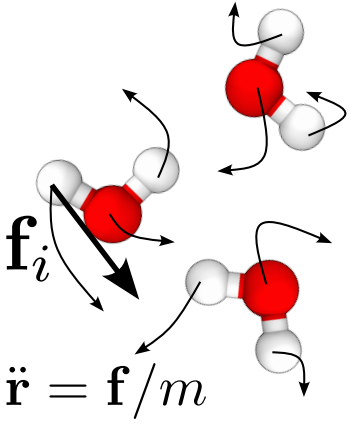

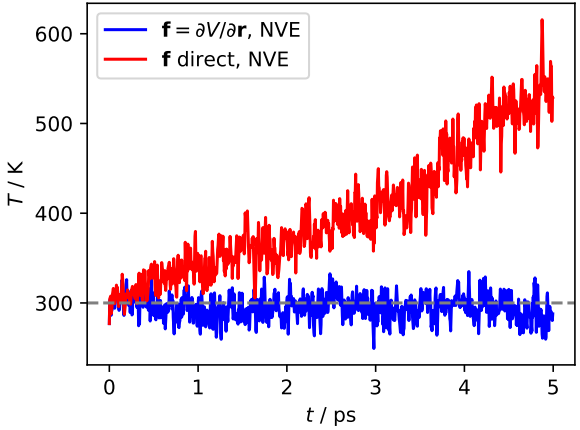


Bigi


Langer

Bigi, Langer, MC, oral @ ICML (2025)

The dark side of the forces

- Many application only need interatomic *forces*, $\mathbf{f}_j = -\partial V\left(\mathbf{r}_1, \ldots \mathbf{r}_N\right) / \partial \mathbf{r}_j$, but learn them as the derivative of a potential energy V
- Data-minded approach: learning $\mathbf{f}_j$ directly is faster
- Bad idea. No energy conservation, unstable sampling, dynamics and geometry optimization

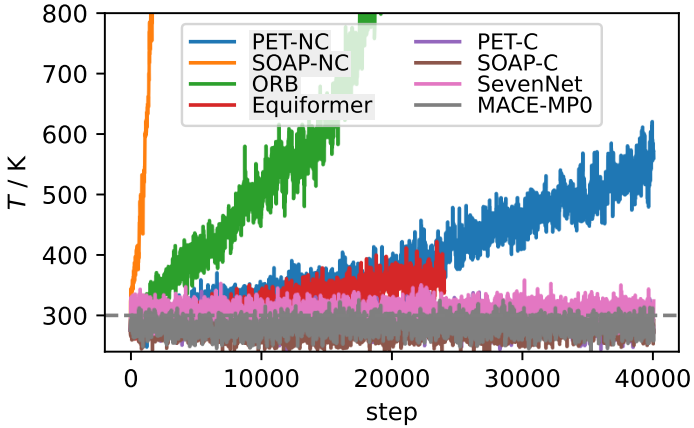


Bigi


Langer

Bigi, Langer, MC, oral @ ICML (2025)

Making good of non-conservative forces

- No easy fix! Thermostatting stabilizes dynamics but leaves serious artifacts
- How about using direct fitting *together* with conservative forces?
  - Conservative fine-tuning: physically-consistent forces at half cost
  - Multiple time stepping: evaluating conservative correction every M steps

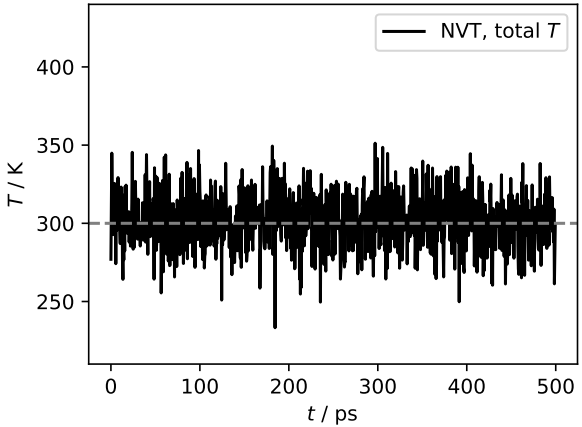


Bigi


Langer

Bigi, Langer, MC, oral @ ICML (2025)

Making good of non-conservative forces

- No easy fix! Thermostatting stabilizes dynamics but leaves serious artifacts
- How about using direct fitting *together* with conservative forces?
  - Conservative fine-tuning: physically-consistent forces at half cost
  - Multiple time stepping: evaluating conservative correction every M steps

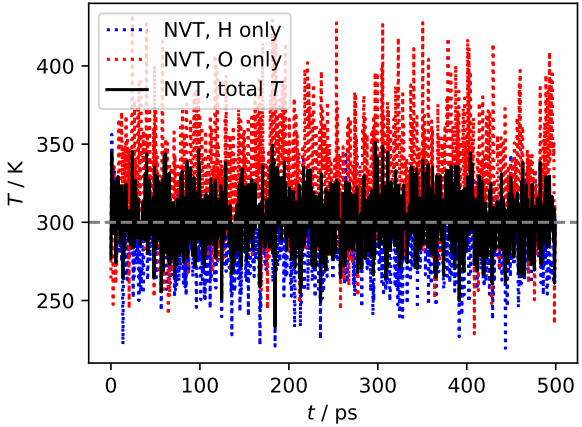


Bigi


Langer

Bigi, Langer, MC, oral @ ICML (2025)

Making good of non-conservative forces

- No easy fix! Thermostatting stabilizes dynamics but leaves serious artifacts
- How about using direct fitting *together* with conservative forces?
  - Conservative fine-tuning: physically-consistent forces at half cost
  - Multiple time stepping: evaluating conservative correction every M steps

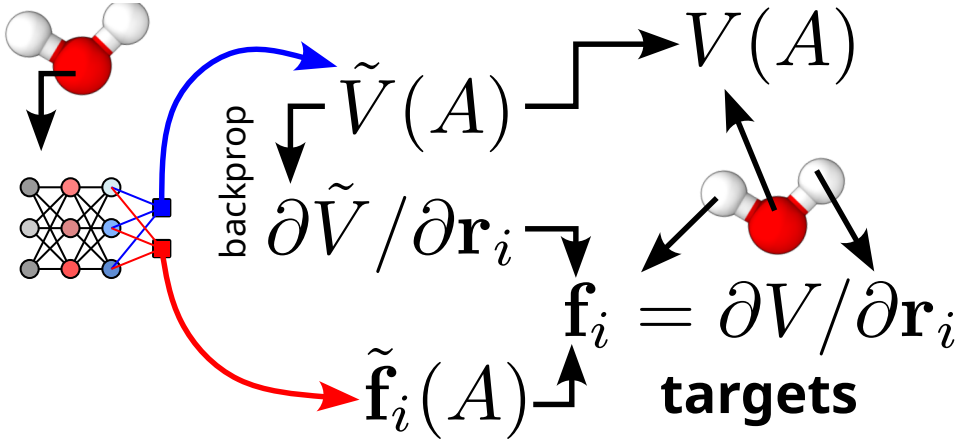


Bigi


Langer

Bigi, Langer, MC, oral @ ICML (2025)

Making good of non-conservative forces

- No easy fix! Thermostatting stabilizes dynamics but leaves serious artifacts
- How about using direct fitting *together* with conservative forces?
  - Conservative fine-tuning: physically-consistent forces at half cost
  - Multiple time stepping: evaluating conservative correction every *M* steps

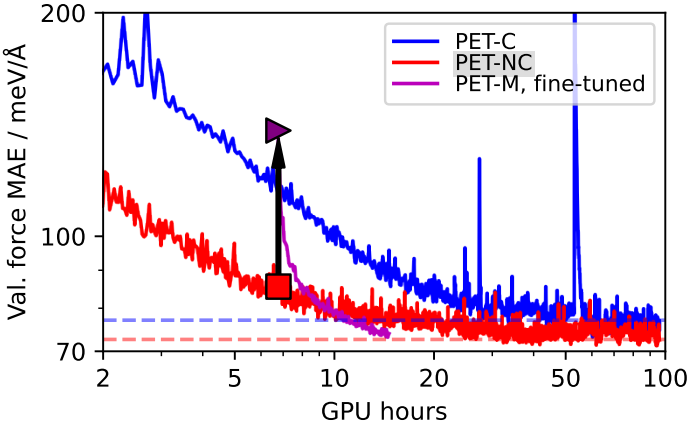


Bigi


Langer

Bigi, Langer, MC, oral @ ICML (2025)

Making good of non-conservative forces

- No easy fix! Thermostatting stabilizes dynamics but leaves serious artifacts
- How about using direct fitting *together* with conservative forces?
  - Conservative fine-tuning: physically-consistent forces at half cost
  - Multiple time stepping: evaluating conservative correction every M steps

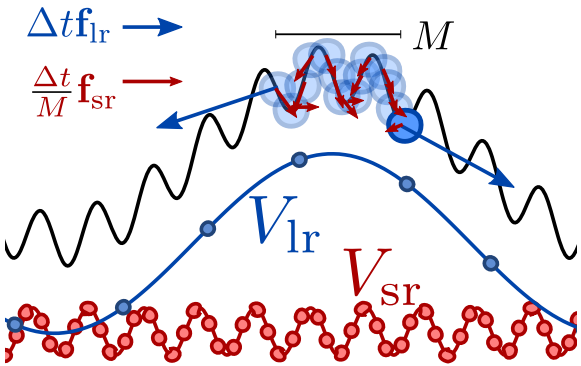

Bigi

Langer

Tuckerman et al., JCP (1992); Bigi, Langer, MC, oral @ ICML (2025)

Making good of non-conservative forces

- No easy fix! Thermostatting stabilizes dynamics but leaves serious artifacts
- How about using direct fitting *together* with conservative forces?
  - Conservative fine-tuning: physically-consistent forces at half cost
  - Multiple time stepping: evaluating conservative correction every M steps

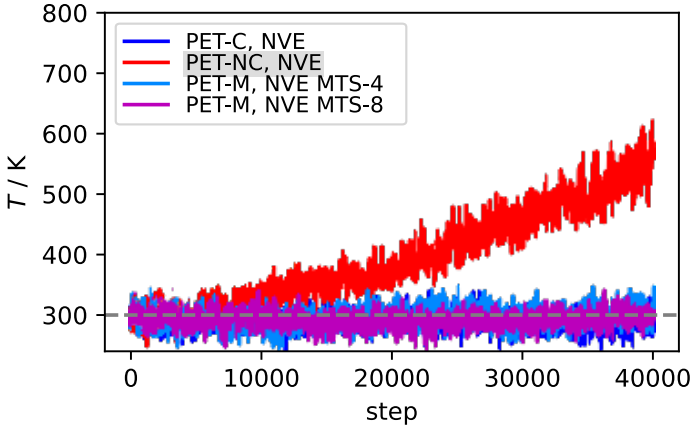


Bigi


Langer

Bigi, Langer, MC, oral @ ICML (2025)

... with the universal potential PET-MAD

- This is a pretty robust setup: works for PET-MAD in the deep extrapolative regime with a simulation of BMIM-Cl at 500K
- See by yourself `https://atomistic-cookbook.org/examples/pet-mad-nc/pet-mad-nc.html`

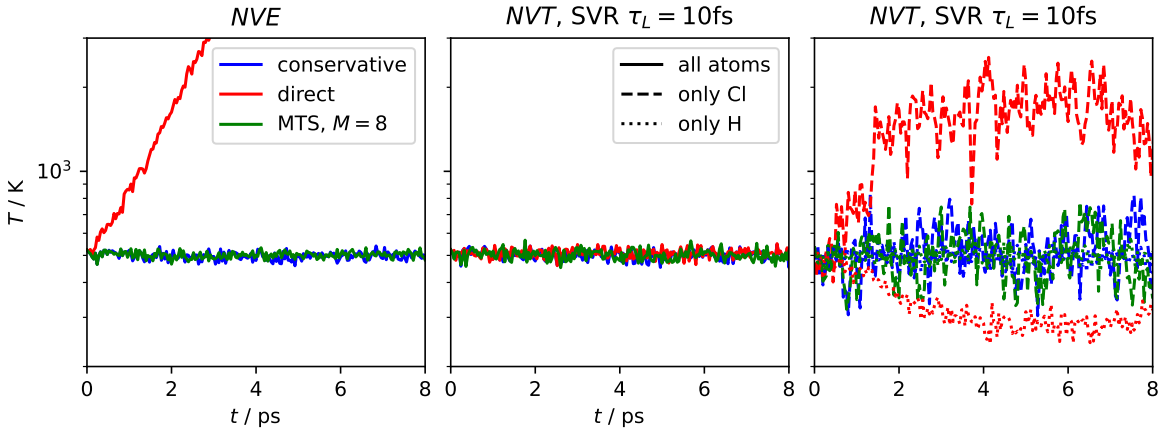

... with the universal potential PET-MAD

- This is a pretty robust setup: works for PET-MAD in the deep extrapolative regime with a simulation of BMIM-Cl at 500K
- See by yourself `https://atomistic-cookbook.org/examples/pet-mad-nc/pet-mad-nc.html`

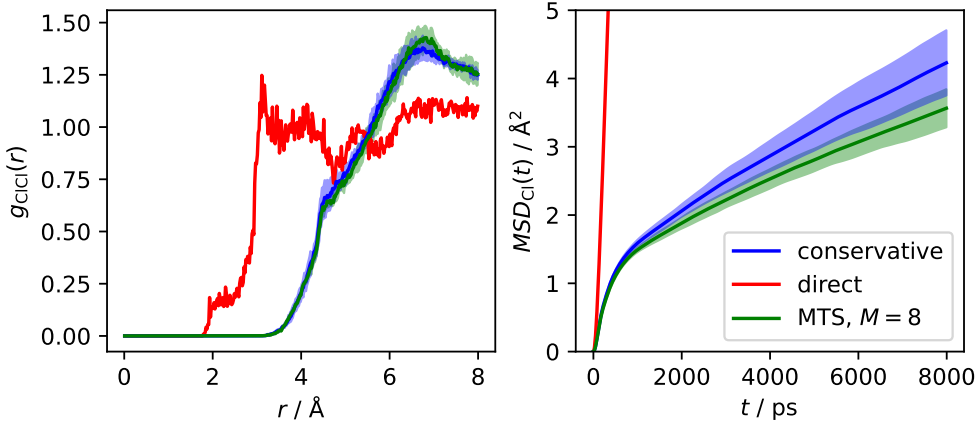

Fast and furious

- If we give up energy conservation, we might at least get more than a 2x speedup. Can we predict $(\mathbf{q}, \mathbf{p})(t + \Delta t)$ starting from $(\mathbf{q}, \mathbf{p})(t)$? With *large* strides?
- Train on 100k MD trajectories generated with PET-MAD → **FlashMD**
- Universal MD with 16-64 fs time steps!

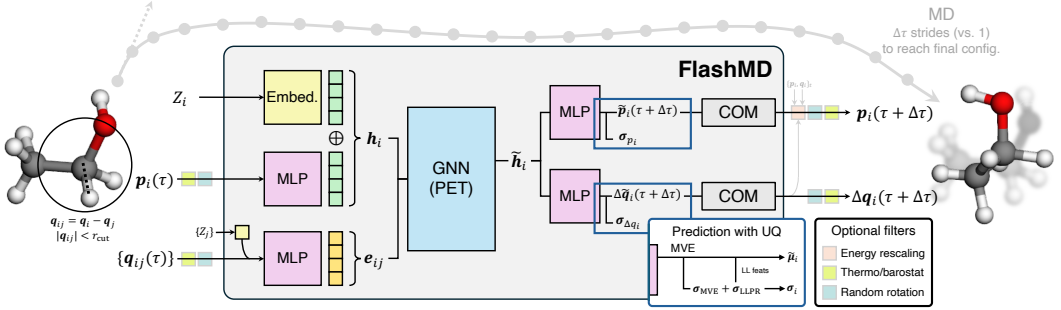

Bigi

Chong

Bigi, Chong, Kristiadi, MC, arxiv:2505.19350

Fast and furious

- If we give up energy conservation, we might at least get more than a 2x speedup. Can we predict $(\mathbf{q}, \mathbf{p})(t + \Delta t)$ starting from $(\mathbf{q}, \mathbf{p})(t)$? With *large* strides?
- Train on 100k MD trajectories generated with PET-MAD $\rightarrow$ **FlashMD**
- Universal MD with 16-64 fs time steps!

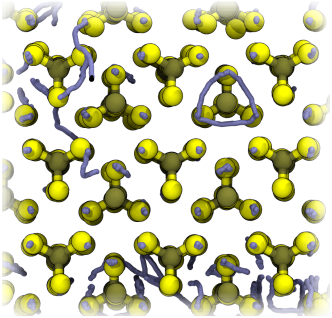

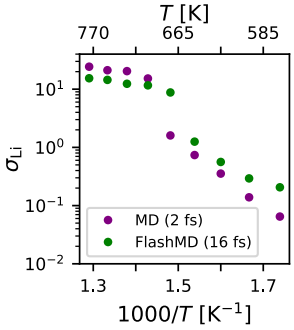

Bigi

Chong

Bigi, Chong, Kristiadi, MC, arxiv:2505.19350

Fast and furious

- If we give up energy conservation, we might at least get more than a 2x speedup. Can we predict $(\mathbf{q}, \mathbf{p})(t + \Delta t)$ starting from $(\mathbf{q}, \mathbf{p})(t)$? With *large* strides?
- Train on 100k MD trajectories generated with PET-MAD $\rightarrow$ **FlashMD**
- Universal MD with 16-64 fs time steps!

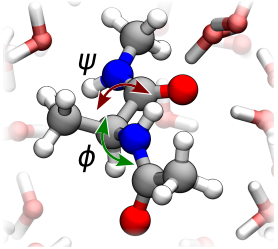

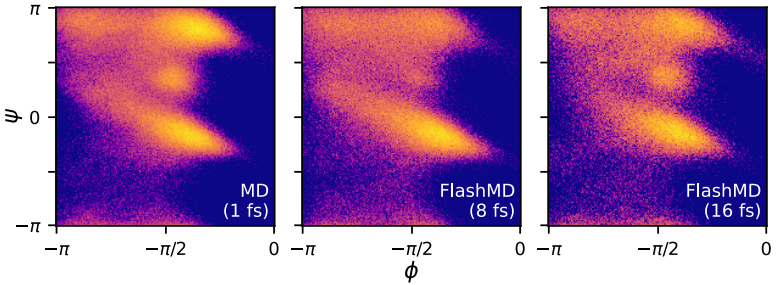

Bigi

Chong

Bigi, Chong, Kristiadi, MC, arxiv:2505.19350

Fast and furious

- If we give up energy conservation, we might at least get more than a 2x speedup. Can we predict $(\mathbf{q}, \mathbf{p})(t + \Delta t)$ starting from $(\mathbf{q}, \mathbf{p})(t)$? With *large* strides?
- Train on 100k MD trajectories generated with PET-MAD $\rightarrow$ **FlashMD**
- Universal MD with 16-64 fs time steps!

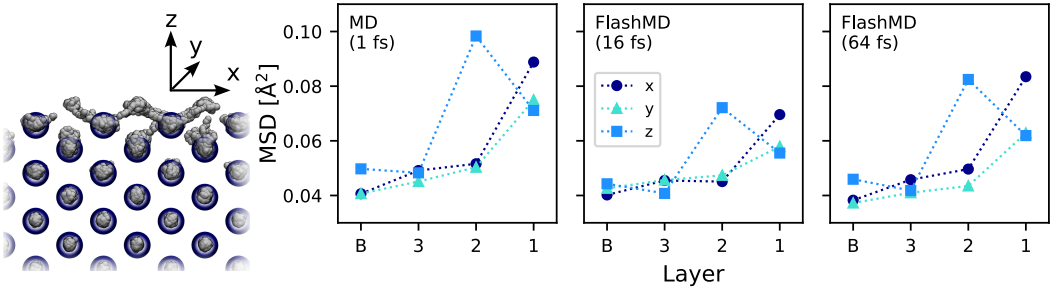

Bigi

Chong

Bigi, Chong, Kristiadi, MC, arxiv:2505.19350

atomistic-cookbook.org/examples/flashmd/flashmd-demo.html

Long-stride trajectories with a universal FlashMD model

Authors:
Michele Ceriotti @ceriottm

This example demonstrates how to run long-stride molecular dynamics using the universal FlashMD model. FlashMD predicts directly positions and momenta of atoms at a later time based on the current positions and momenta. It is trained on MD trajectories obtained with the PET-MAD universal potential. You can read more about the model and its limitations in this preprint.

```
# sphinx_gallery_thumbnail_path = '../../examples/flashmd/flashmd-scheme.png'
```

Start by importing the required libraries. You will need the PET-MAD potential, as well as FlashMD and a recent version of i-PI.

```
pip install pet-mad flashmd ipi
```

```
import chemiscope
from flashmd import get_universal_model
from flashmd.ipi import get_npt_stepper, get_nvt_stepper
from ipi.utils.parsing import read_output, read_trajectory
from ipi.utils.scripting import InteractiveSimulation
from pet_mad.calculator import PETMADCalculator
```

A rough schematic of the architecture of FlashMD is shown below. Each model is trained for a specific

Thank you!

cosmo.epfl.ch @labcosmo

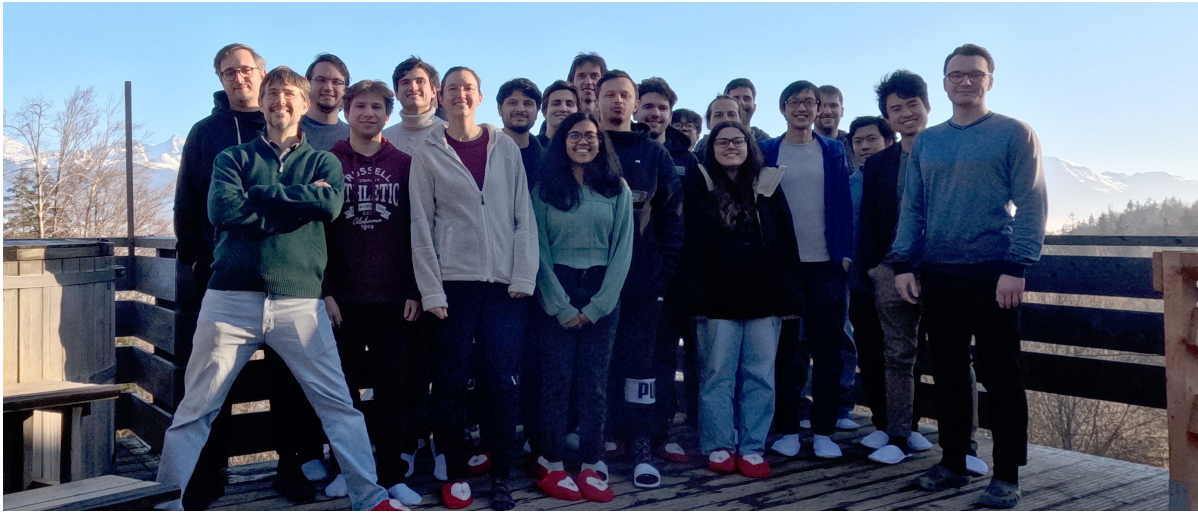

Thank you!


CONGRATULATIONS !
YOU ARE NOW AT UHURU PEAK
5895 M A.M.S.L.
TANZANIA
AFRICA'S HIGHEST POINT
FREE STANDING MOUNTAIN
RGP

Last-layer ensembles are size-consistent

- Atomistic potentials are often sums of atom-centered predictions $y(A) = \sum_{i \in A} y(A_i)$
- How do atomic errors propagate to structure errors?
- Mixed behavior for a water model in the NVT ensemble
- Well-calibrated force prediction

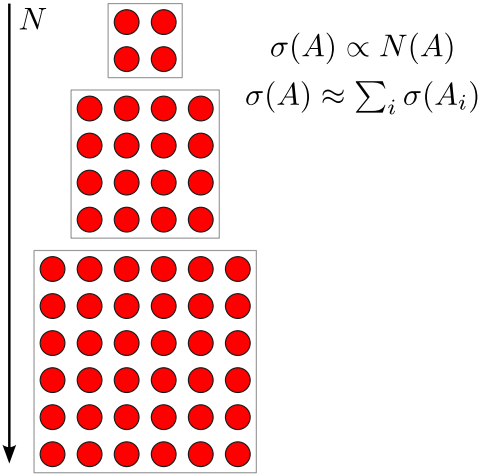

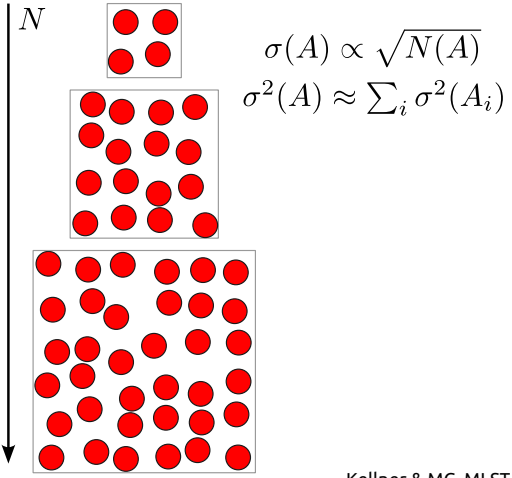

Kellner & MC, MLST (2024)

Last-layer ensembles are size-consistent

- Atomistic potentials are often sums of atom-centered predictions $y\,(A) = \sum_{i \in A} y\,(A_i)$
- How do atomic errors propagate to structure errors?
- Mixed behavior for a water model in the NVT ensemble
- Well-calibrated force prediction

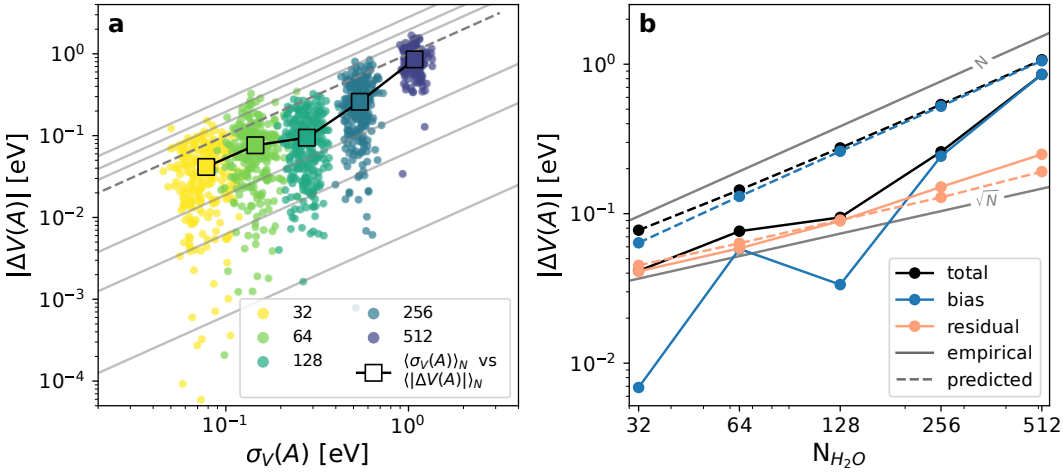

Kellner & MC, MLST (2024)

Last-layer ensembles are size-consistent

- Atomistic potentials are often sums of atom-centered predictions $y\,(A) = \sum_{i \in A} y\,(A_i)$
- How do atomic errors propagate to structure errors?
- Mixed behavior for a water model in the NVT ensemble
- Well-calibrated force prediction

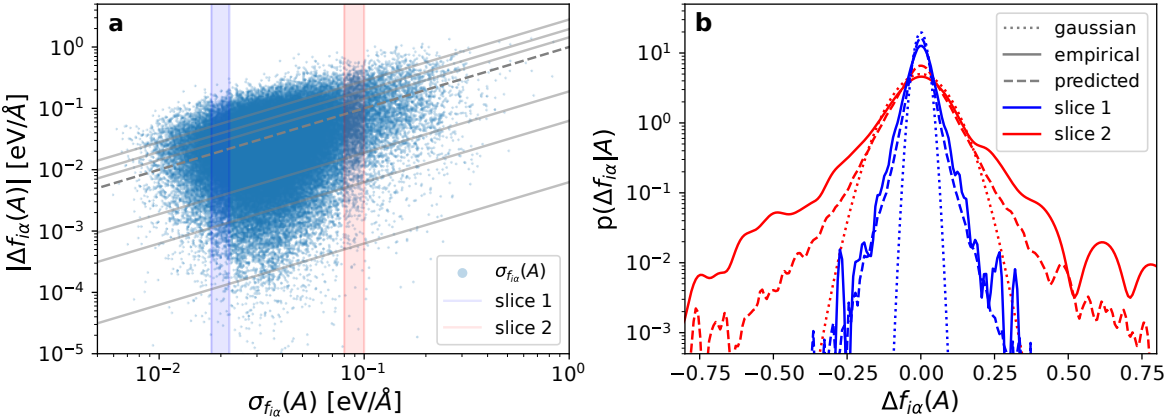

Kellner & MC, MLST (2024)

Last-layer ensembles extrapolate well

- Accuracy degrades for out-of-sample predictions, but is good enough to select structures for active learning
- Incremental training improves UQ calibration as well as prediction accuracy

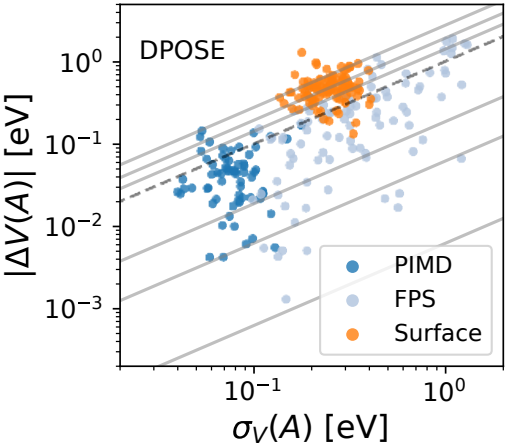

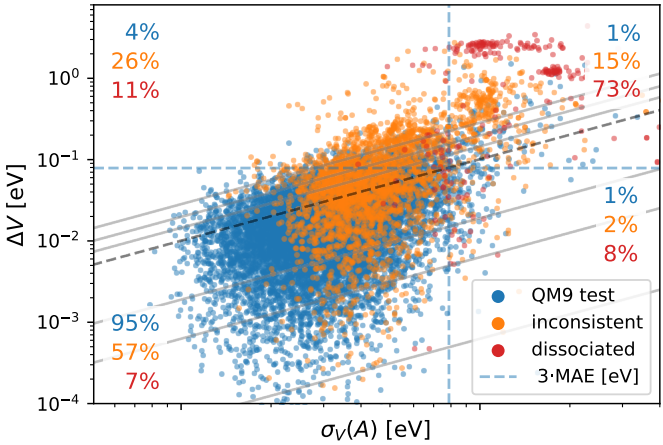

Kellner & MC, MLST (2024)

Last-layer ensembles extrapolate well

- Accuracy degrades for out-of-sample predictions, but is good enough to select structures for active learning
- Incremental training improves UQ calibration as well as prediction accuracy

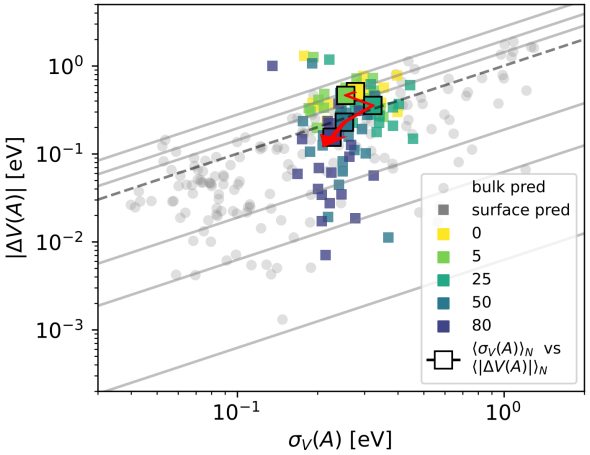

Kellner & MC, MLST (2024)

UQ for acid-base equilibria

- Energetics described with a DFTB baseline and a MLP correction. Accelerated by multiple time stepping
- PIGLET thermostatting & ring-polymer contraction for the quantum sampling, Plumed-driven metadynamics applied to the centroid
- Metadynamics sampling of the dissociation of CH_3SO_2OH.
- **UQ for the free-energy profile!**

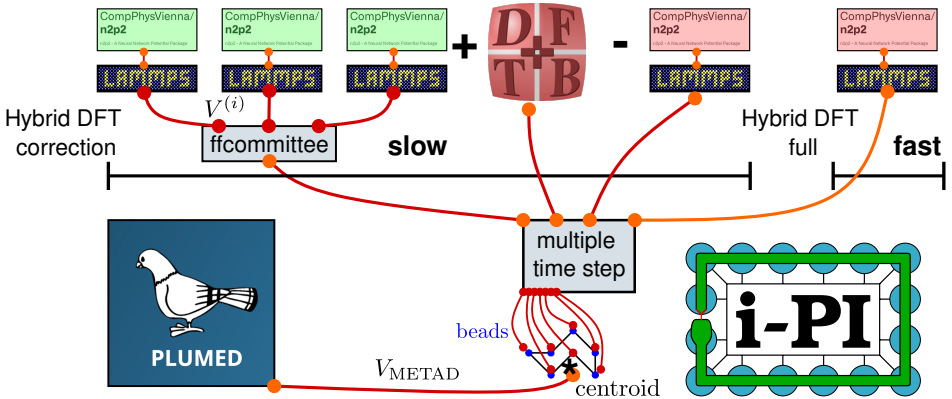

Kapil, Behler, MC, JCP (2016); Zamani et al., Adv. Mater. (2020); Rossi et al., JCTC (2020)

UQ for acid-base equilibria

- Energetics described with a DFTB baseline and a MLP correction. Accelerated by multiple time stepping
- PIGLET thermostatting & ring-polymer contraction for the quantum sampling, Plumed-driven metadynamics applied to the centroid
- Metadynamics sampling of the dissociation of CH_3SO_2OH.
- **UQ for the free-energy profile!**

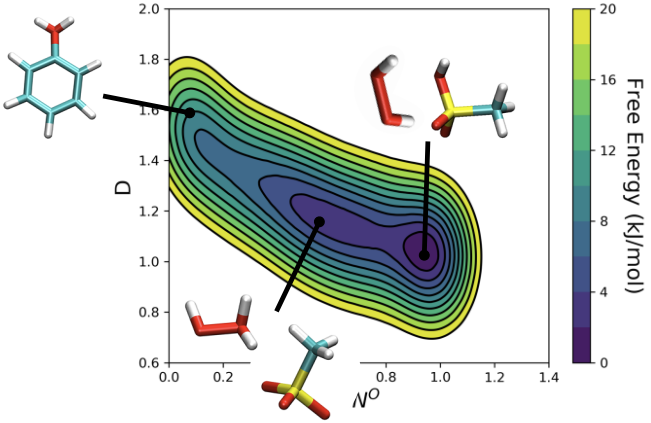

Kapil, Behler, MC, JCP (2016); Zamani et al., Adv. Mater. (2020); Rossi et al., JCTC (2020)

UQ for acid-base equilibria

- Energetics described with a DFTB baseline and a MLP correction. Accelerated by multiple time stepping
- PIGLET thermostatting & ring-polymer contraction for the quantum sampling, Plumed-driven metadynamics applied to the centroid
- Metadynamics sampling of the dissociation of CH_3SO_2OH.
- **UQ for the free-energy profile!**

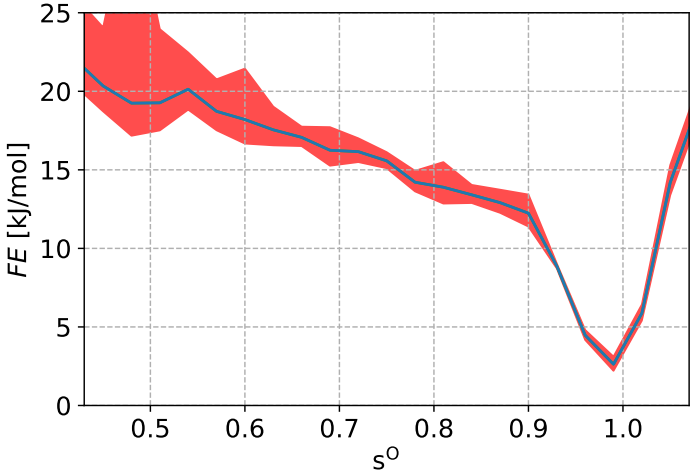

Imbalzano et al., JCP (2021)

Going MAD

- Training set generation based on Maximum Atomic Diversity.
  Internally consistent DFT details - aim for DFT accuracy with < 100k structures
- Diverse data space, covering more ground than much larger datasets
- The latent variables provide a general framework for data analytics

Subset name	Description	# structures	# atoms
MC3D	Bulk crystals from the Materials Cloud 3D crystals	33596	738484
MC3D-rattled	Rattled analogs of the original MC3D crystals, with Gaussian noise added to all atomic positions	30044	599675
MC3D-random	Artificial structures from MC3D with randomized atomic species from 85 elements	2800	25095
MC3D-surface	Surface slabs generated from MC3D by cleaving along random low-index crystallographic planes	5589	205185
MC3D-cluster	Nanoclusters (2-8 atoms) cut from MC3D and MC3D-rattled environments	9071	44829
MC2D	2D crystals from the Materials Cloud	2676	43225
SHIFTML-molcrys	Curated SHIFTML molecular crystals from the Cambridge Structural Database	8578	852044
SHIFTML-molfrags	Neutral molecular fragments from SHIFTML	3241	72120

Going MAD

- Training set generation based on Maximum Atomic Diversity.
  Internally consistent DFT details - aim for DFT accuracy with < 100k structures
- Diverse data space, covering more ground than much larger datasets
- The latent variables provide a general framework for data analytics

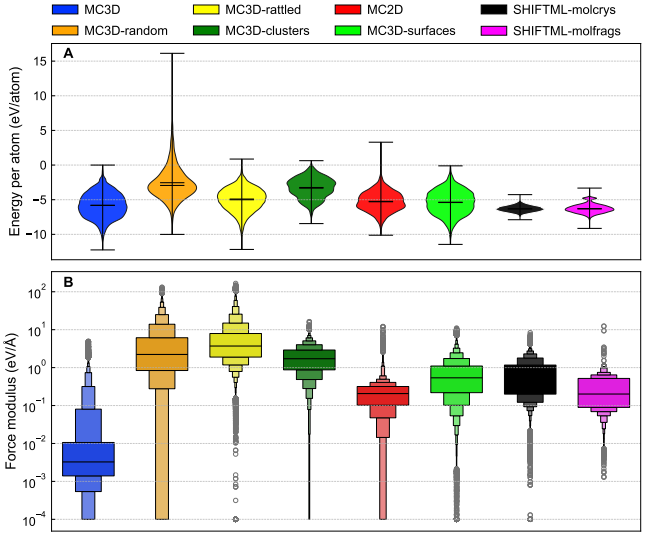

Going MAD

- Training set generation based on Maximum Atomic Diversity.
  Internally consistent DFT details - aim for DFT accuracy with < 100k structures
- Diverse data space, covering more ground than much larger datasets
- The latent variables provide a general framework for data analytics

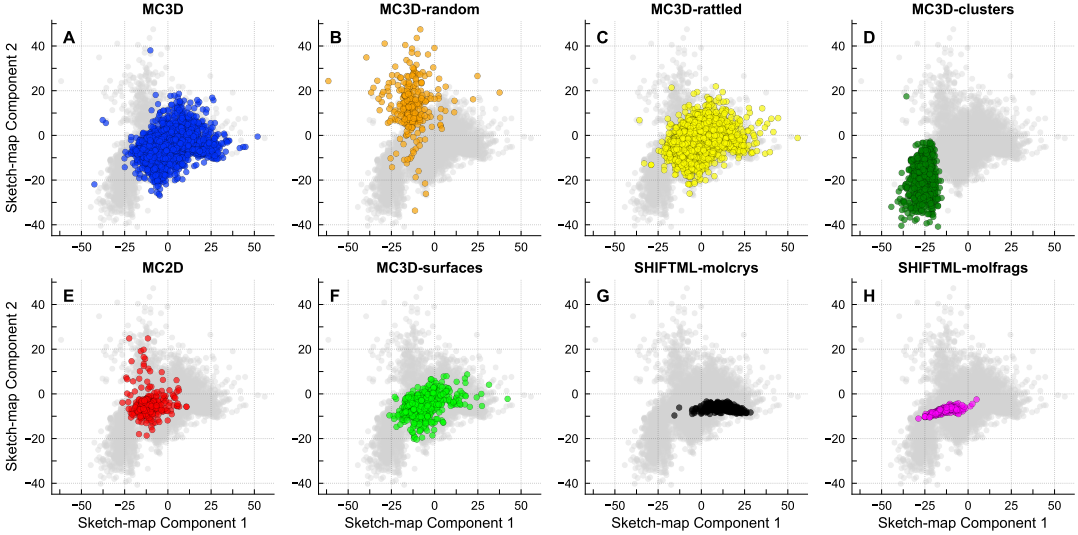

Going MAD

- Training set generation based on Maximum Atomic Diversity.
  Internally consistent DFT details - aim for DFT accuracy with < 100k structures
- Diverse data space, covering more ground than much larger datasets
- The latent variables provide a general framework for data analytics

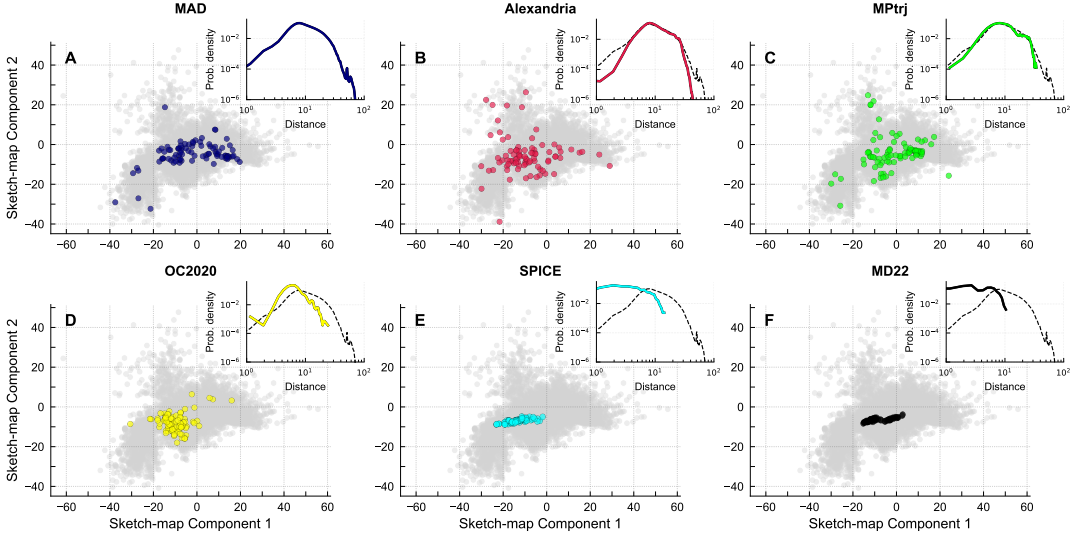

Going MAD

- Training set generation based on Maximum Atomic Diversity.
  Internally consistent DFT details - aim for DFT accuracy with < 100k structures
- Diverse data space, covering more ground than much larger datasets
- The latent variables provide a general framework for data analytics

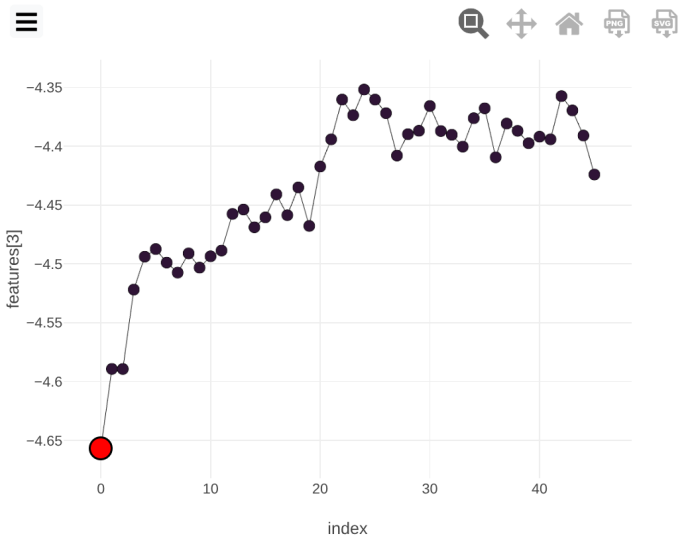

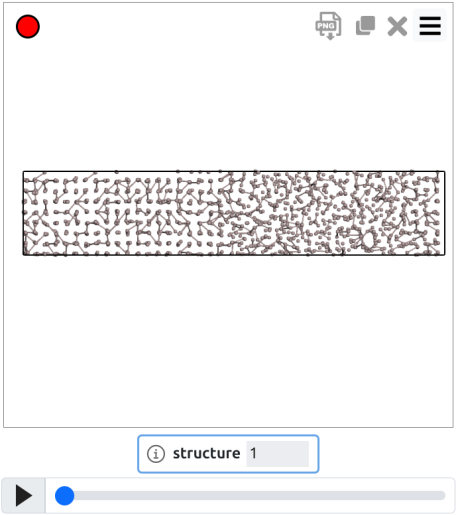

Going MAD

- Training set generation based on Maximum Atomic Diversity.
  Internally consistent DFT details - aim for DFT accuracy with < 100k structures
- Diverse data space, covering more ground than much larger datasets
- The latent variables provide a general framework for data analytics

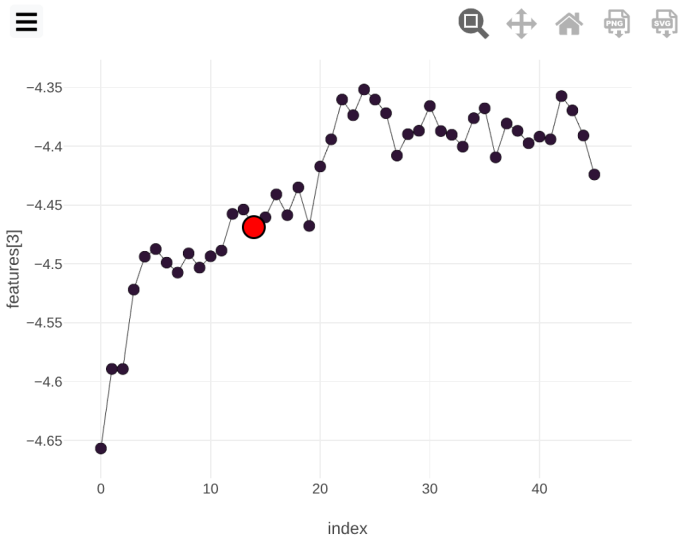

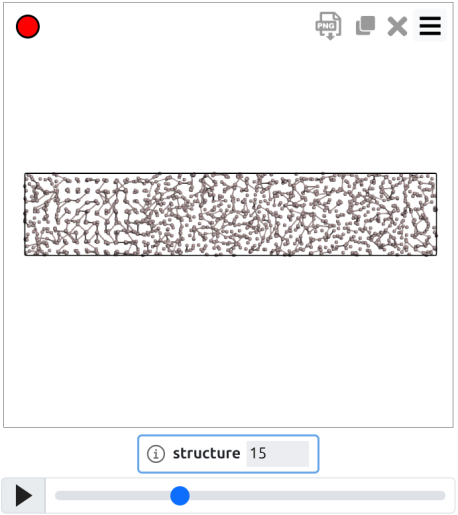

Going MAD

- Training set generation based on Maximum Atomic Diversity.
  Internally consistent DFT details - aim for DFT accuracy with < 100k structures
- Diverse data space, covering more ground than much larger datasets
- The latent variables provide a general framework for data analytics

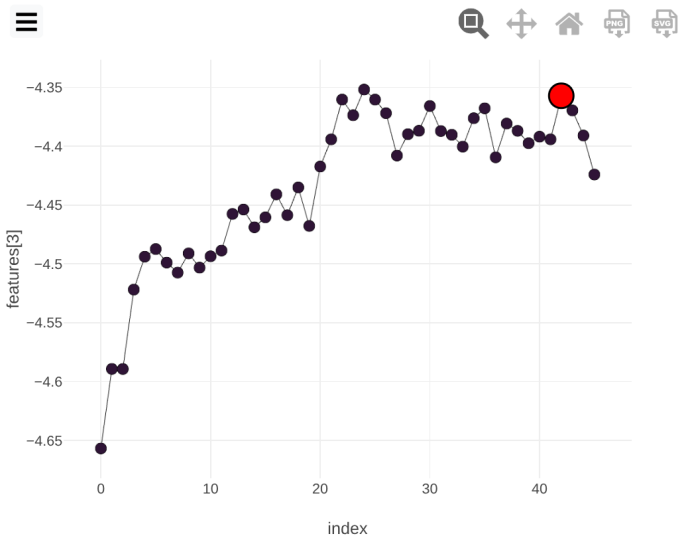

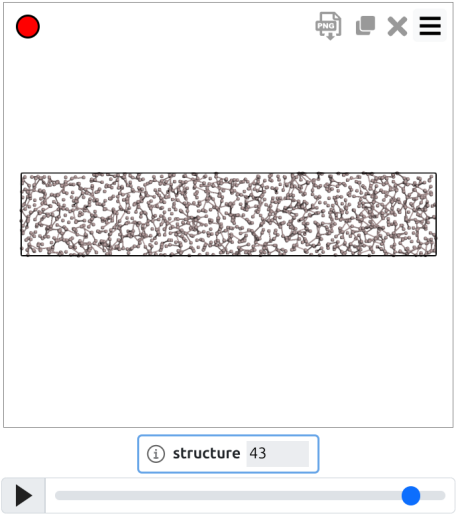

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

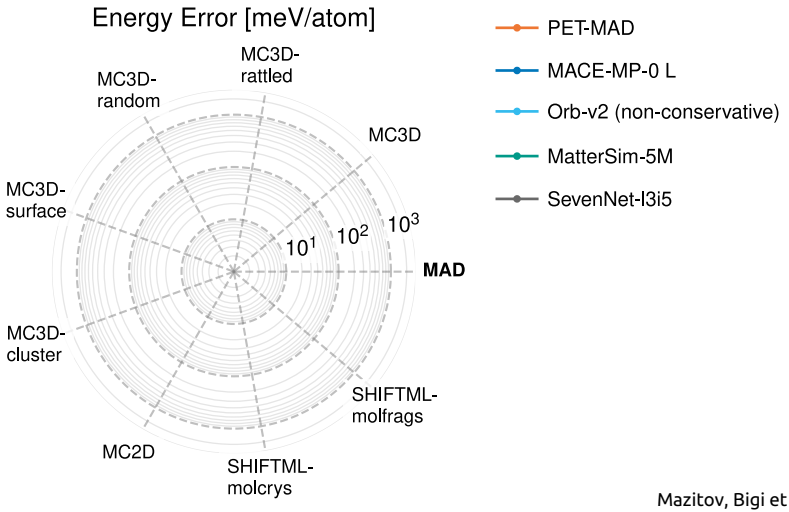


Mazitov


Bigi

Mazitov, Bigi et al., arXiv:2503.14118;
`https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html`

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

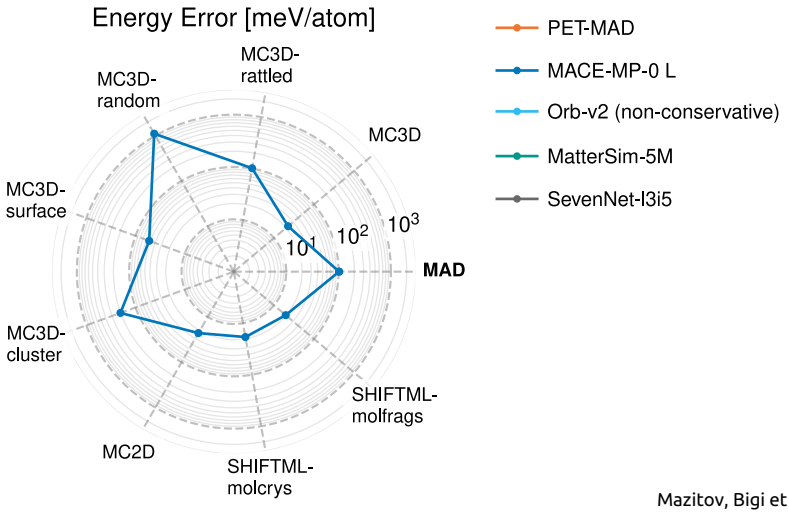


Mazitov


Bigi

Mazitov, Bigi et al., arXiv:2503.14118;
`https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html`

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

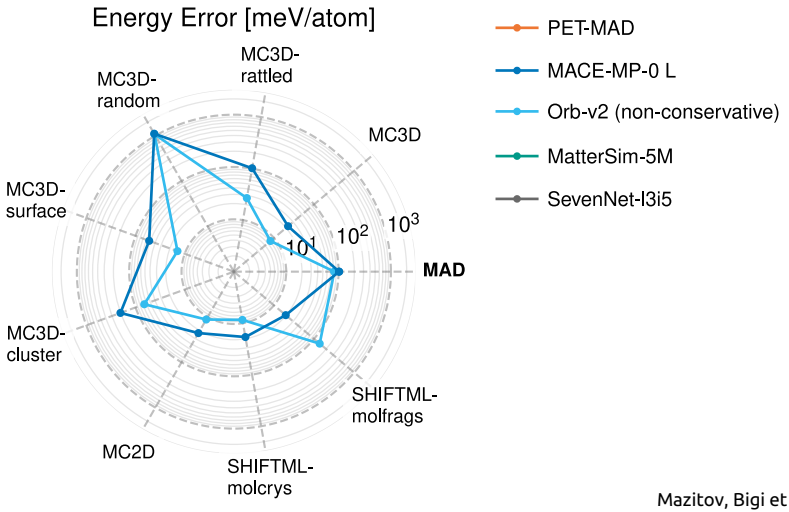

Mazitov

Bigi

Mazitov, Bigi et al., arXiv:2503.14118;
https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

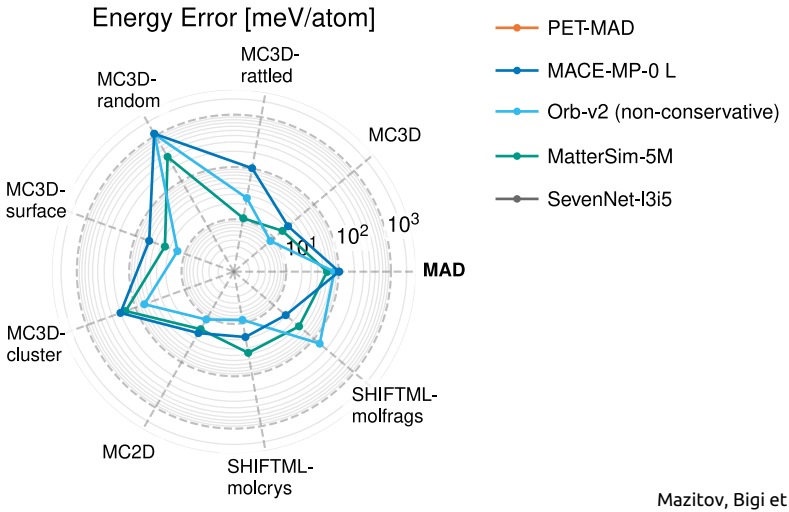

Mazitov

Bigi

Mazitov, Bigi et al., arXiv:2503.14118;
https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

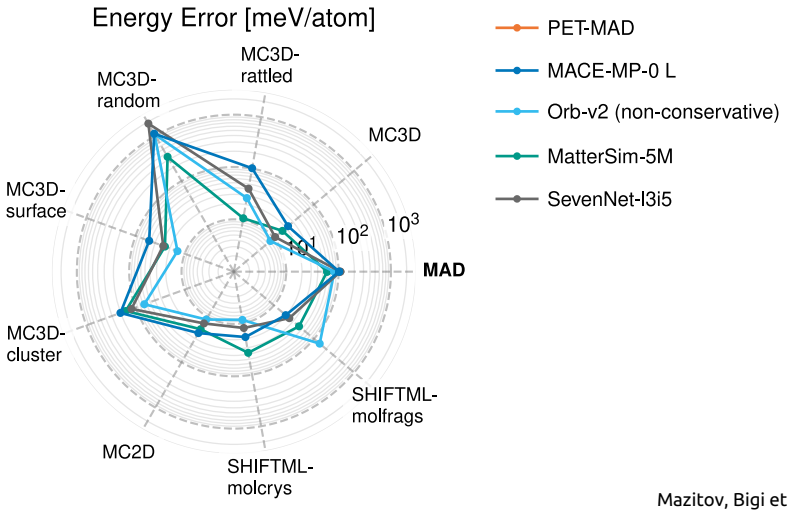


Mazitov


Bigi

Mazitov, Bigi et al., arXiv:2503.14118;
`https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html`

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

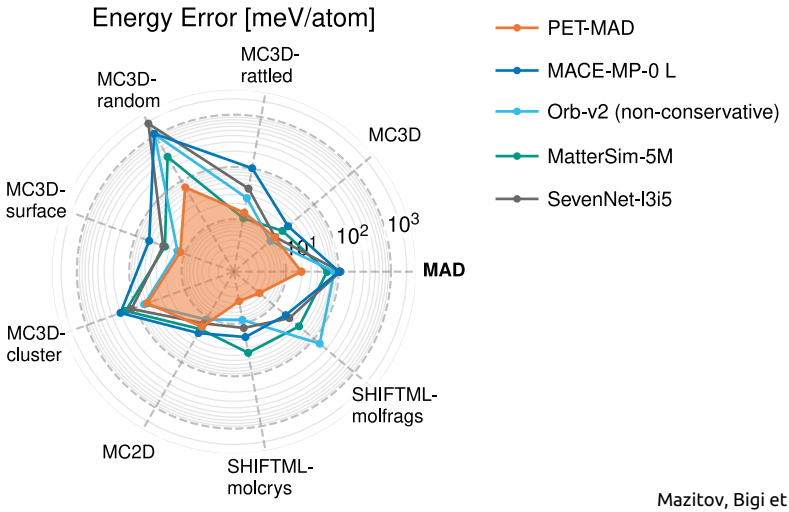


Mazitov


Bigi

Mazitov, Bigi et al., arXiv:2503.14118;
`https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html`

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

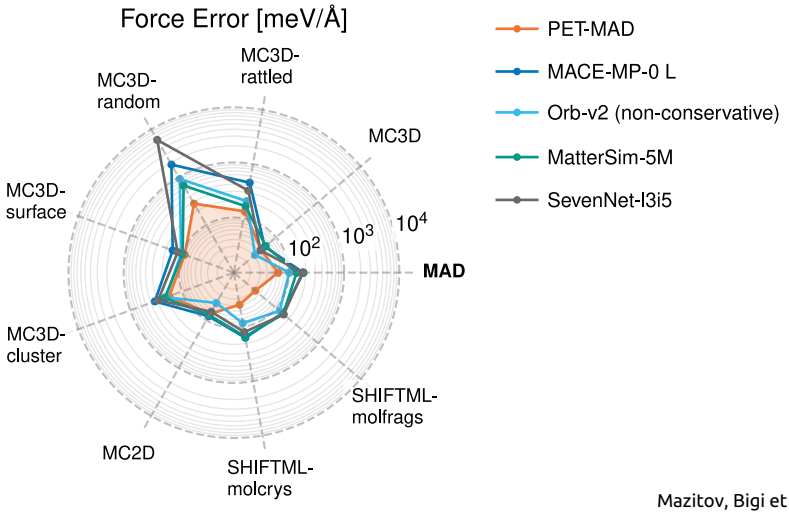

Mazitov

Bigi

Mazitov, Bigi et al., arXiv:2503.14118;
https://atomistic-cookbook.org/examples/pet-mad/pet-mad.html

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

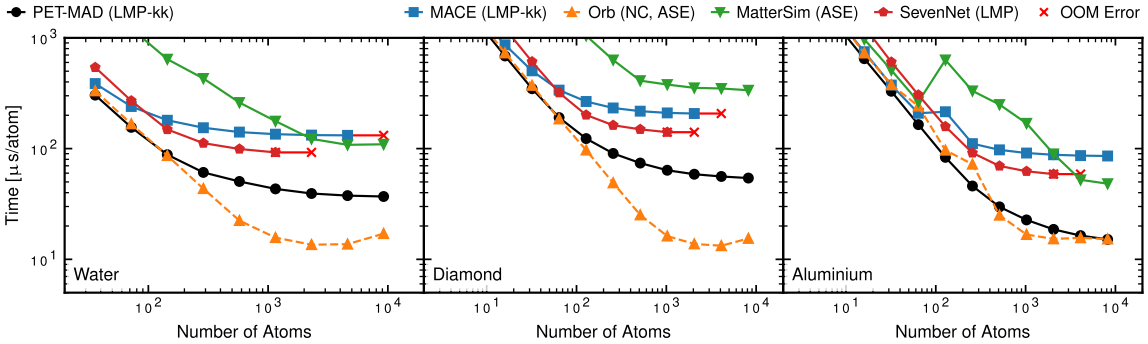

Mazitov, Bigi et al., arXiv:2503.14118;

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

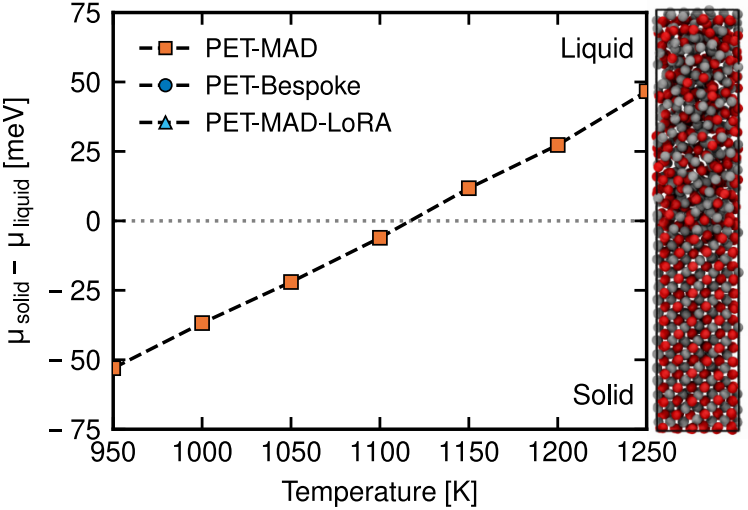


Kellner


Mazitov


Bigi

Mazitov, Bigi et al., arXiv:2503.14118;

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

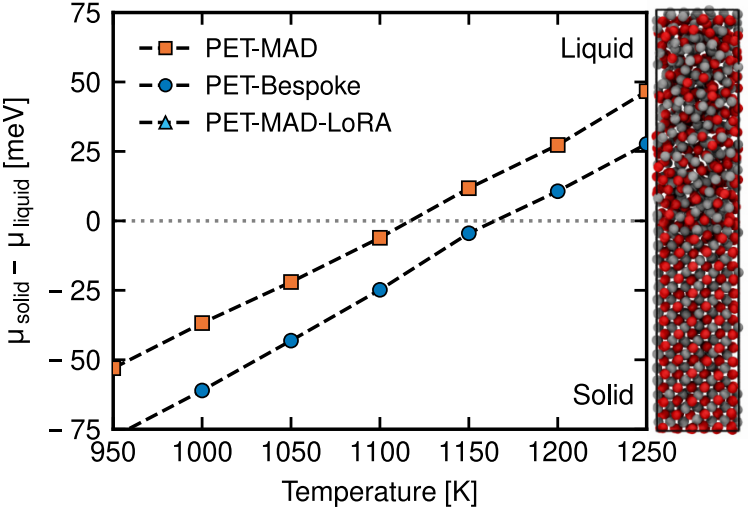


Kellner


Mazitov


Bigi

Mazitov, Bigi et al., arXiv:2503.14118;

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

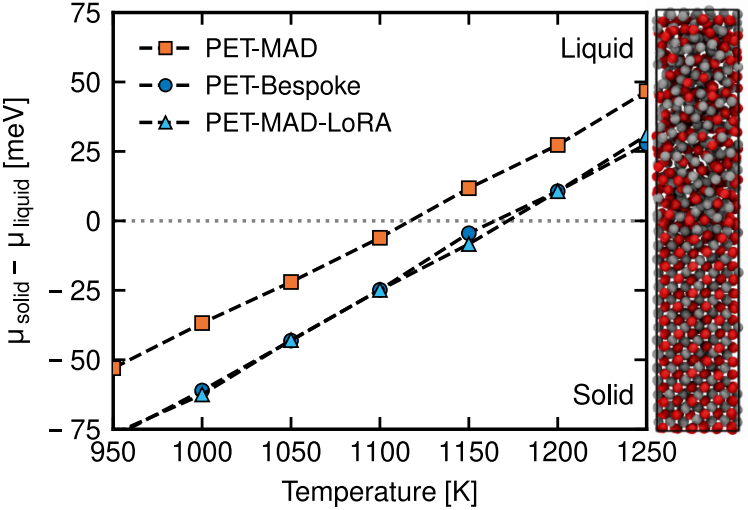


Kellner


Mazitov


Bigi

Mazitov, Bigi et al., arXiv:2503.14118;

PET-MAD - Accuracy across materials (and molecular) space

- **Consistent** benchmarking shows competitive performance with 100x less data
- Fast, unconstrained architecture
- Prediction of complex properties, simple fine-tuning, model errors with shallow ensembles

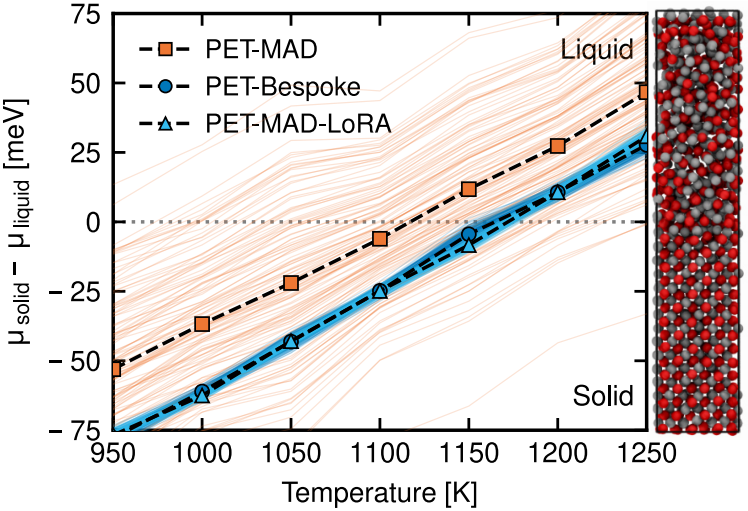


Kellner


Mazitov


Bigi

Kellner, Ceriotti, MLST (2025); Mazitov, Bigi et al., arXiv:2503.14118;