

Basic principles of application of machine learning in quantum chemistry

Jan Hermann

FU Berlin, Department of Mathematics, Berlin, Germany https://jan.hermann.name

VŠCHT, 2 Jun 2022

Computational modeling of molecules and materials: What and Why?

molecules and materials

lab



Computational modeling of molecules and materials: How?

• Schrödinger equation – exact in principle, impossible to solve exactly in practice



Atoms and electrons: Thermodynamics and electronic structure

- Molecules made of nuclei (\mathbf{R}_I) and electrons (\mathbf{r}_i)
- Total molecular wave function

 $\psi(\mathbf{r}_1,\ldots,\mathbf{r}_N,\mathbf{R}_1,\ldots,\mathbf{R}_M)$

• Born–Oppenheimer approximation: separation of nuclear and electronic motion

$$(\mathbf{R}_{1},...,\mathbf{R}_{N}) \rightarrow \hat{H}_{elec} \rightarrow \psi_{n}(\mathbf{r}_{1},...,\mathbf{r}_{N}), E_{n}$$

$$E_{n}(\mathbf{R}_{1},...,\mathbf{R}_{M}) \rightarrow \hat{H}_{nucl}$$

$$potential \ energy \ surface$$

$$(\mathbf{R}_{1},...,\mathbf{R}_{M}) \rightarrow \hat{H}_{nucl}$$



Electronic structure problem



- Ab-initio: exact **1**, approximate **2** (*coupled clusters, quantum Monte Carlo*)
- Semi-empirical: approximate **1**, exact **2** (*density functional theory, tight-binding, Hubbard model*)
- Machine learning (force fields, density functionals, wave function ansatzes)

Machine learning: Data and objectives

- Human learning: World model from observation and intervention
- Machine learning: Algorithms improve from empirical evidence
- Data: Formalized empirical evidence, subset of "ground truth"
- Objective: Measure of quality of a model
- Generalization: Does the model work on new empirical evidence?

Machine learning flavors



Machine learning: Glorified fitting?

"With four parameters I can fit an elephant, and with five I can make him wiggle his trunk." —J. von Neumann



| neural networks | number of parameters |
|--------------------|--------------------------|
| toy models | ~ 10 ³ |
| serious models | >10 ⁵ |
| AlphaFold 2 | ~108 |
| GPT-3 | ~1011 |

Mayer et al. Drawing an elephant with four complex parameters. AJP 78, 648 (2010)

Bias-variance tradeoff and overfitting





• Increasing number of parameters eventually leads to overfitting

@daniela_witten

Regularization and high dimensionality



@daniela_witten



- Regularization—picking the most smooth fit
- Balestriero et al.: Learning in High Dimension Always Amounts to Extrapolation

Machine learning and "double descent"



Nakkiran et al. Deep Double Descent (ICRL, 2019)

Gaussian process regression

 $f(x) = \sin(2\pi x)$







- Straightforward generalization to multiple dimensions and arbitrary *descriptors*
- Needs good descriptors
- "Just" a linear combination limited flexibility
- Limited to small datasets
- Largely replaced with neural networks

From Gaussian processes to neural networks

$$y = \sum_{k} W_{k} e^{-\frac{(x-x_{k})^{2}}{\sigma^{2}}}$$
$$= \sum_{k} W_{k} e^{-(\sigma^{-1}x-\sigma^{-1}x_{k})^{2}}$$
$$= \sum_{k} W_{k} e^{-(W_{k}'x-b_{k})^{2}}$$
$$= \sum_{k} W_{k} g(W_{k}'x-b_{k})$$

$$\mathbf{y} = \mathbf{W}g(\mathbf{W}'\mathbf{x} - \mathbf{b})$$



$$\mathbf{y} = \mathbf{W}g(\mathbf{W}'g(\mathbf{W}''\mathbf{x}-\mathbf{b}')-\mathbf{b})$$





Neural networks and deep learning

$$\mathbf{y} = \mathbf{W}g(\mathbf{W}'g(\mathbf{W}''\xi(\mathbf{x})-\mathbf{b}')-\mathbf{b}) \qquad \xi: \checkmark \to \mathbb{R}^n$$

- Deep neural networks are "universal approximators" – good descriptors unnecessary
- Cannot be optimized analytically gradient descent
- Fast evaluation of gradients via backpropagation
- Stochastic gradient descent from minibatches (implicit regularization)

$$\boldsymbol{\theta} := \boldsymbol{\theta} - \eta \nabla_{\boldsymbol{\theta}} \mathscr{L} \big(\boldsymbol{\theta}, \{ \mathbf{x}_i \}_{i \in B_n} \big)$$



Some deep learning successes

COMPUTER SCIENCE

A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play

David Silver^{1,2*†}, Thomas Hubert^{1*}, Julian Schrittwieser^{1*}, Ioannis Antonoglou¹, Matthew Lal¹, Arthur Guez¹, Marc Lanctot¹, Laurent Sifre¹, Dharshan Kumaran¹, Thore Graepel¹, Timothy Lillicrap¹, Karen Simonyan¹, Demis Hassabis¹[†]

The game of chess is the longest-studied domain in the history of artificial intelligence. The strongest programs are based on a combination of sophisticated search techniques, domain-specific adaptations, and handcrafted evaluation functions that have been refined by human experts over several decades. By contrast, the AlphaGo Zero program recently achieved superhuman performance in the game of Go by reinforcement learning from self-play. In this paper, we generalize this approach into a single AlphaZero algorithm that can achieve superhuman performance in many challenging games. Starting from random play and given no domain knowledge except the game rules, AlphaZero convincingly defeated a world champion program in the games of chess and shogi (Japanese chess), as well as Go. programmers, combined with a high-performance alpha-beta search that expands a vast search tree by using a large number of clever heuristics and domain-specific adaptations. In (10) we describe these augmentations, focusing on the 2016 Top Chess Engine Championship (TCEC) season 9 world champion Stockfish (11); other strong chess programs, including Deep Blue, use very similar architectures (1, 12).

In terms of game tree complexity, shogi is a substantially harder game than chess (13, 14): It is played on a larger board with a wider variety of pieces; any captured opponent piece switches sides and may subsequently be dropped anywhere on the board. The strongest shogi programs, such as the 2017 Computer Shogi Association (CSA) world champion Elmo, have only recently defeated human champions (15). These programs use an algorithm similar to those used by computer chess programs, again based on a highly optimized alpha-beta search engine with many domain-specific adaptations.

AlphaZero replaces the handcrafted knowledge and domain-specific augmentations used in traditional game-playing programs with deep

Article

Highly accurate protein structure prediction with AlphaFold

https://doi.org/10.1038/s41586-021-03819-2 Received: 11 May 2021 Accepted: 12 July 2021 Published online: 15 July 2021 Open access Check for updates John Jumper¹⁴, Richard Evans¹⁴, Alexander Pritzel¹⁴, Tim Green¹⁴, Michael Figurnov¹⁴, Olaf Ronneberger¹⁴, Kathryn Tunyasuvunakool¹⁴, Russ Bates¹⁴, Augustin Žídek¹⁴, Anna Potapenko¹⁴, Alex Bridgland¹⁴, Clemens Meyer¹⁴, Simon A. A. Kohl¹⁴, Andrew J. Ballard¹⁴, Andrew Cowie¹⁴, Bernardino Romera-Paredes¹⁴, Stanislav Nikolov¹⁴, Rishub Jain¹⁴, Jonas Adler¹, Trevor Back¹, Stig Petersen¹, David Reiman¹, Ellen Clancy¹, Michal Zielinski¹, Martin Steinegger²³, Michalina Pacholska¹, Tamas Berghammer¹, Sebastian Bodenstein¹, David Silver¹, Oriol Vinyals¹, Andrew W. Senior¹, Koray Kavukcuoglu¹, Pushmeet Kohli¹ & Demis Hassabis¹⁴

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort¹⁻⁴, the structures of around 100,000 unique proteins have been determined⁵, but this represents a small fraction of the billions of known protein sequences⁶⁷. Structural coverage is bottlenecked by the months to years of painstaking effort required to determine a single protein structure. Accurate computational approaches are needed to address this gap and to enable large-scale structural bioinformatics. Predicting the

RESEARCH

QUANTUM CHEMISTRY

Pushing the frontiers of density functionals by solving the fractional electron problem

James Kirkpatrick¹*†, Brendan McMorrow¹†, David H. P. Turban¹†, Alexander L. Gaunt¹†, James S. Spencer¹, Alexander G. D. G. Matthews¹, Annette Obika¹, Louis Thiry², Meire Fortunato¹, David Pfau¹, Lara Román Castellanos¹, Stig Petersen¹, Alexander W. R. Nelson¹, Pushmeet Kohli¹, Paula Mori-Sánchez³, Demis Hassabis¹, Aron J. Cohen^{1.4}*

Density functional theory describes matter at the quantum level, be suffer from systematic errors that arise from the violation of mathe functional. We overcame this fundamental limitation by training a n data and on fictitious systems with fractional charge and spin. The (DeepMind 21), correctly describes typical examples of artificial cha correlation and performs better than traditional functionals on thor atoms and molecules. DM21 accurately models complex systems su charged DNA base pairs, and diradical transition states. More crucis methodology relies on data and constraints, which are continually in viable pathway toward the exact universal functional. code [PySCF (19)]. The functional was evaluated by integrating local energies computed by a multilayer perceptron (MLP), which took as input both local and nonlocal features of the occupied Kohn-Sham (KS) orbitals, and can be described as a local range-separated hybrid. To train the functional, the sum of two objective functions was used: a regression loss for learning the exchange-correlation energy itself and a gradient regularization term that ensured that the functional field (SCE) cal-



Magnetic control of tokamak plasmas through deep reinforcement learning

https://doi.org/10.1038/s41586-021-04301-9 Received: 14 July 2021 Accepted: 1 December 2021 Published online: 16 February 2022

Jonas Degrave¹³, Federico Felici^{2,39}, Jonas Buchli^{13,9}, Michael Neunert¹³, Brendan Tracey^{13,9}, Francesco Carpanese^{12,3}, Timo Ewalds¹³, Roland Hafner¹³, Abbas Abdolmalek Diego de las Casas¹, Craig Donner¹, Leslie Fritz¹, Cristian Galperti², Andrea Huber¹, James Keeling¹, Maria Tsimpoukelli¹, Jackie Kay¹, Antoine Merle², Jean-Marc Moret², Seb Noury¹, Federico Pesamosca², David Pfau¹, Olivier Sauter², Cristian Sommariva², Stefano Coda², Basil Duval², Ambrogio Fasol², Pushmeet Kohli¹, Koray Kavukcuoglu¹, Demis Hassabis¹ & Martin Riedmiller¹³

Control policy

- 2

Onon

Article



Nuclear fusion using magnetic confinement, in particular in the tokamak configuration, is a promising path towards sustainable energy. A core challenge is to shape and maintain a high-temperature plasma within the tokamak vessel. This requires high-dimensional, high-frequency, closed-loop control using magnetic actuator coils, further complicated by the diverse requirements across a wide range of plasma configurations. In this work, we introduce a previously undescribed

plasma configurations. In this work, we introduce a previously undescribed architecture for tokamak magnetic controller design that autonomously learns to command the full set of control coils. This architecture meets control objectives specified at a high level, at the same time satisfying physical and operational constraints. This approach has unprecedented flexibility and generality in problem specification and yields a notable reduction in design effort to produce new plasma

Symmetry and equivariance













Symmetry and equivariance

- Symmetry group set (group) of transformations *g* leaving an object *x* unchanged (invariant)
- Equivariant map (function) map *f* from a symmetric object to another object that preserves symmetry

$$f(T_g x) = T_g f(x)$$

• If objects are (lists of) numbers (scalars, vectors, tensors), the transformations are encoded in irreducible representations

$$\mathbf{f}(\mathbf{D}_{\mu}(g)\mathbf{x}) = \mathbf{D}_{\nu}(g)\mathbf{f}(\mathbf{x})$$

• Invariance – special case of equivariance for the trivial representation

$$\mathbf{f}(\mathbf{D}_{\mu}(g)\mathbf{x}) = \mathbf{f}(\mathbf{x})$$

Symmetries and equivariance in machine learning

- Without explicit treatment, ML models are not equivariant
- Poor man's equivariance data augmentation
- Invariance easy through invariant descriptors

$$f_{\boldsymbol{\theta}}\big(\boldsymbol{\xi}(\mathbf{D}_{\boldsymbol{\mu}}(g)\mathbf{x})\big) = f_{\boldsymbol{\theta}}\big(\boldsymbol{\xi}(\mathbf{x})\big)$$

- General equivariance must be built into the ML model
- Examples of symmetry preserving operations

$$\mathbf{v} \to f(|\mathbf{v}|)\hat{\mathbf{v}} \qquad x_i \to g(x_i) \sum_j h(x_j)$$

Machine learning in practice

- Parameters vs. hyperparameters
- Dataset splitting



Unke et al. Machine learning force fields. Chem. Rev. 121, 10142 (2021)



Machine-learning force fields

• Circumvent high cost of electronic-structure methods via a surrogate force field model for molecular dynamics



Unke et al. Machine learning force fields. Chem. Rev. 121, 10142 (2021)

Machine-learning force fields

- Given a few thousand reference structures, energies, and forces, train $E_{\theta}(\{Z_i, \mathbf{R}_i\})$
- Training set needs to cover what's sampled during MD
 - interpolation vs. extrapolation?
 - active learning, uncertainty, ensembles
- Cost of ab-initio MD vs dataset preparation, training, and ML MD



Unke et al. Machine learning force fields. Chem. Rev. 121, 10142 (2021)

$\mathbf{x}_i = (Z_i, \mathbf{R}_i)$

Machine-learning force-field models

• Traditional many-body expansion

 $E(\{\mathbf{x}_i\}) = \sum_i E^{(1)}(\mathbf{x}_i) + \frac{1}{2} \sum_{ij} E^{(2)}(\{\mathbf{x}_i, \mathbf{x}_j\}) + \frac{1}{6} \sum_{ijk} E^{(3)}(\{\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k\}) + \dots$

• Semi-local atomic energy expansion $E(\{\mathbf{x}_i\}) = \sum_i \varepsilon(\mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}) = \sum_i f_{\theta}(\boldsymbol{\xi}(\mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}))$



• Many-body expansion of local environment

 $\boldsymbol{\xi}(\mathbf{x}_{i}, \{\mathbf{x}_{j}\}_{j \in \mathcal{N}_{i}}) = \left[\boldsymbol{\xi}^{(1)}(\mathbf{x}_{i}), \{\boldsymbol{\xi}^{(2)}(\mathbf{x}_{i}, \mathbf{x}_{j})\}_{j \in \mathcal{N}_{i}}, \{\boldsymbol{\xi}^{(3)}(\mathbf{x}_{i}, \mathbf{x}_{j}, \mathbf{x}_{k})\}_{j,k \in \mathcal{N}_{i}}, \dots\right]$

• Local environment descriptors from message passing

$$\boldsymbol{\xi}_{i}^{(n+1)} = g_{\boldsymbol{\theta}} \left(\boldsymbol{\xi}_{i}^{(n)}, \{ \boldsymbol{\xi}_{j}^{(n)} \}_{j \in \mathcal{N}_{i}} \right)$$



 $\mathbf{x}_i = (Z_i, \mathbf{R}_i)$

Radial and angular basis for atomic environments

the only rotationally invariant one-atom descriptor



Zoo of atomic descriptors

• Good news – convergence towards a single framework



Musil et al. Physics-Inspired Structural Representations. Chem. Rev. 121, 9759 (2021)

Simplified classification of ML force fields

| | Fixed descriptor | Message-passing descriptor |
|---------------------|--|--|
| Gaussian process | SOAP, FCHL, | |
| Neural network | Behler–Parrinello , DeepMD, ANI, | SchNet, PhysNet, NequIP, PaiNN, DimeNet, NewtonNet, |

Example 1: SOAP–GAP

Article Origins of structural and electronic transitions in disordered silicon

| https://doi.org/10.1038/s41586-020-03072-z | Volker L. Deringer ^{1⊠} , Noam Bernstein², Gábor Csányi³, Chiheb Ben Mahmoud ^{4,5} , | |
|--|---|--|
| Received: 14 December 2019 | Michele Ceriotti ^{4,5} , Mark Wilson ⁶ , David A. Drabold ⁷ & Stephen R. Elliott ^{8,9} | |
| Accepted: 12 November 2020 | | |
| Published online: 6 January 2021 | Structurally disordered materials pose fundamental questions ¹⁻⁴ , including how different disordered phases ('polyamorphs') can coexist and transform from one phase to another ⁵⁻⁹ . Amorphous silicon has been extensively studied; it forms a | |
| Check for updates | | |

- Molecular dynamics: 100,000 atoms, 10³-nm cell, 10 ps
- Training: PW91, <200 atoms per uc, 2,500 structures, 170,000 atomic environments



Progress of compression run

Example 2: Behler– Parrinello

Ab initio thermodynamics of liquid and solid water

Bingqing Cheng^{a,1}, Edgar A. Engel^a, Jörg Behler^{b,c}, Christoph Dellago^d, and Michele Ceriotti^a

^aLaboratory of Computational Science and Modeling, Institute of Materials, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland; ^bUniversität Göttingen, Institut für Physikalische Chemie, Theoretische Chemie, 37077 Göttingen, Germany; ^cInternational Center for Advanced Studies of Energy Conversion, Universität Göttingen, 37073 Göttingen, Germany; and ^dFaculty of Physics, University of Vienna, 1090 Vienna, Austria

Edited by Pablo G. Debenedetti, Princeton University, Princeton, NJ, and approved December 3, 2018 (received for review September 4, 2018)

Thermodynamic properties of liquid water as well as hexagonal (Ih) and cubic (Ic) ice are predicted based on density functional theory at the hybrid-functional level, rigorously taking into account quantum nuclear motion, anharmonic fluctuations, and machine-learning (ML) techniques to avoid the prohibitively large computational expenses otherwise incurred by extensively sampling phase space by using first-principles methods. In particular, we use sophisticated thermodynamic integration (TI)

- Molecular dynamics: 800 molecules, nuclear quantum effects
- Training: revPBE0-D3, 64 molecules, 1,600 structures



Example 3: SchNet

Roaming leads to amino acid photodamage: A deep learning study of tyrosine

Julia Westermayr,¹ Michael Gastegger,² Dora Vörös,³ Lisa Panzenboeck,^{3,4} Florian Joerg,^{3,5} Leticia González,^{3,6} and Philipp Marquetand^{3,6,7,a)}

¹⁾University of Warwick, Department of Chemistry, Gibbet Hill Rd, Coventry, CV4 7AL,

UK

²⁾Technical University of Berlin, Machine Learning Group, 10587 Berlin, Germany

³⁾University of Vienna, Faculty of Chemistry, Institute of Theoretical Chemistry, Währinger Str. 17, 1000 Vienna, Austria.

⁴⁾Present address: University of Vienna, Faculty of Chemistry, Department (1090 Vienna, Austria.

⁵⁾Present address: University of Vienna, Faculty of Chemistry, Institute of Co 1000 Vienna, Austria nature chemistry

- Molecular dynamics: 1000
 1ps trajectories, 29 electronic states
- Training: CASPT2(12e,11o), 15,000 structures







Electronic Schrödinger equation

$$\hat{H}\psi(\mathbf{r}_1, \dots, \mathbf{r}_N) = E\psi(\mathbf{r}_1, \dots, \mathbf{r}_N)$$

$$\hat{H} := \sum_{i=1}^N \left(-\frac{1}{2} \nabla_{\mathbf{r}_i}^2 - \sum_I \frac{Z_I}{|\mathbf{r}_i - \mathbf{R}_I|} \right) + \sum_{i < j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}$$

$$\hat{h}_i$$

- Electron coordinates: $(\mathbf{r}_1, ..., \mathbf{r}_N) \equiv \mathbf{r}$
- Molecule specified by atom charges Z_I and coordinates \mathbf{R}_I
- Solution: eigenstates ψ_n and their energies E_n , including the ground state ψ_0, E_0
- Variational principle: $E_0 = \min_{\psi} \int d\mathbf{r} \, \psi(\mathbf{r}) \hat{H} \psi(\mathbf{r})$
- Analytically solvable only for hydrogen atom, but many approximate numerical methods from quantum chemistry



Machine-learned density functionals

• Density-functional theory (DFT) exact in principle, but only approximate functionals known in practice

$$\hat{H} := \sum_{i} \left(\hat{h}_{i} + v_{\text{eff}}[n](\mathbf{r}_{i}) \right), \quad n(\mathbf{r}) = \sum_{i} \phi_{i}(\mathbf{r})^{2}$$
$$v_{\text{eff}}[n](\mathbf{r}) = \int d\mathbf{r}' \left(\frac{n(\mathbf{r})}{|\mathbf{r} - \mathbf{r}'|} + \frac{\delta E_{\text{xc}}[n]}{\delta n(\mathbf{r})} \right), \quad E_{\text{xc}}[n] \approx \int d\mathbf{r} \, \epsilon_{\text{xc}}(n(r))$$

- Given reference energies, learn functional? 🗙
- DFT calculation involves minimization w.r.t. energy
- Given reference energies and densities, learn functional $\overline{\mathbf{V}}$

End-to-end learning with autodifferentiation



Li et al. Kohn-Sham equations as regularizer. PRL 126, 036401 (2021)

Incorporating physics improves generalization



• Only two points needed to fit dissociation of H_2

Li et al. Kohn-Sham equations as regularizer. PRL 126, 036401 (2021)

Example: DM21

QUANTUM CHEMISTRY

Pushing the frontiers of density functionals by solving the fractional electron problem

James Kirkpatrick¹*⁺, Brendan McMorrow¹⁺, David H. P. Turban¹⁺, Alexander L. Gaunt¹⁺, James S. Spencer¹, Alexander G. D. G. Matthews¹, Annette Obika¹, Louis Thiry², Meire Fortunato¹, David Pfau¹, Lara Román Castellanos¹, Stig Petersen¹, Alexander W. R. Nelson¹, Pushmeet Kohli¹, Paula Mori-Sánchez³, Demis Hassabis¹, Aron J. Cohen^{1,4}* code [PySCF (19)]. The functional was evaluated by integrating local energies computed by a multilayer perceptron (MLP), which took as input both local and nonlocal features of the occupied Kohn-Sham (KS) orbitals, and can be described as a local range-separated hybrid. To train the functional, the sum of two objective functions was used: a regression loss for learning the exchange-correlation energy itself and a gradient regularization term that ensured that the functional derivatives

- DM21: Range-separated meta-GGA hybrid, comparable in accuracy and training to state-of-the-art functionals
- Exact constraints through data augmentation





Machine-learned wave functions

Chem. Listy 112, 640-647 (2018)

JE STROJOVÉ UČENÍ BUDOUCNOSTÍ TEORETICKÉ CHEMIE?

KAREL BERKA^a, ŠTĚPÁN SRŠEŇ^b a PETR SLAVÍČEK^{b,c}

Numerické algoritmy strojového učení ale mohou změnit přímo i *ab initio* kvantovou chemii. Přesné výpočty jsou zde komplikovány korelacemi mezi pohyby jednotlivých částic, což vede k exponenciálně se zvyšujícím nárokům s rostoucí velikostí systému. Techniky strojového učení aplikované na samotnou vlnovou funkci mohou být v některých případech schopny výrazně výpočty urychlit⁴⁸.

⁴⁸Carleo & Troyer, *Science* **355**, 602 (2017)

Variational quantum Monte Carlo

Schrödinger equation **Basis states**, wave function **Monte Carlo** integration Local energy Variational principle **Gradient descent** Markov-chain sampling

 $\hat{H}|\psi\rangle = E|\psi\rangle$ $|\psi\rangle = \sum_{\mathbf{x}} |\mathbf{x}\rangle \langle \mathbf{x} |\psi\rangle = \sum_{\mathbf{x}} \psi(\mathbf{x}) |\psi\rangle$ $E[\psi] = \frac{\langle \psi | H | \psi \rangle}{\langle \psi | \psi \rangle} = \mathbb{E}_{\mathbf{x} \sim |\psi(\mathbf{x})|^2} [E_{\text{loc}}(\mathbf{x})]$ $E_{\text{loc}}(\mathbf{x}) = \sum_{\mathbf{x}'} \langle \mathbf{x} | \hat{H} | \mathbf{x}' \rangle \frac{\psi(\mathbf{x}')}{\psi(\mathbf{x})}$ (zero variance) $E = \min_{\psi} E[\psi] \le \min_{\theta} E[\psi_{\theta}]$ $\nabla_{\boldsymbol{\theta}} E[\psi_{\boldsymbol{\theta}}] = \mathbb{E}_{\mathbf{x}} \left[\left(E_{\text{loc}}(\mathbf{x}) - \mathbb{E}_{\mathbf{x}'}[E_{\text{loc}}(\mathbf{x}')] \right) \nabla_{\boldsymbol{\theta}} \ln |\psi_{\boldsymbol{\theta}}(\mathbf{x})| \right]$

 $\mathbf{x}' \sim \exp(-\|\mathbf{x} - \mathbf{x}'\|/\tau), \quad P(\mathbf{x} := \mathbf{x}') = |\psi(\mathbf{x}')|^2 / |\psi(\mathbf{x})|^2$

Deep variational quantum Monte Carlo

```
\nabla_{\theta} E
def fit wf(
    hamil: Hamiltonian[X],
    wf: Callable[[X], Psi],
    sampler: Iterator[X],
    opt: torch.optim.Optimizer,
    steps: int,
                                               \mathbf{x} \sim |\boldsymbol{\psi}|^2
):
    for _ in zip(steps):
         x = next(sampler)
         psi = wf(x)
         E_loc = hamil.local_energy(x, wf)
         loss = (
              (E_loc - E_loc.mean()).detach() * psi.log()
         ).mean()
         loss.backward()
         opt.step()
         opt.zero_grad()
```

First and second quantization

- Quantum states of many-particle systems must be ullet(anti)symmetric for bosons (fermions) with respect to exchange of particles
- Two ways to treat the (anti)symmetry ullet

First quantization

(one-particle discrete basis: $\mathbf{x} \rightarrow k$) Second quantization

$$|\psi\rangle = \sum_{n_1 \cdots n_K} \psi(n_1, \dots, n_K) |n_1 \cdots n_K\rangle$$

occupation numbers n_i

 $\mathbf{n} = (0, 0, 0, 2, 0, 1)$

$$N = 3, K = 5$$

$$\mathbf{x} = (2.7, 3.1, 4.9)$$

$$\mathbf{X} = (2.7, 5.1, 4.9)$$

Electronic Schrödinger equation (first quantization)



- Electron coordinates, $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_N)$, 3N-dimensional space
- (Asymmetric) basis states are $|\mathbf{r}\rangle$, antisymmetric wave function $\psi(\mathbf{r})$



- Molecule specified by atom charges Z_I and coordinates \mathbf{R}_I
- Local energy needs evaluation of Laplacian

$$E_{\text{loc}}(\mathbf{r}) = -\frac{1}{2} \frac{\Delta \psi(\mathbf{r})}{\psi(\mathbf{r})} + \frac{1}{2} \sum_{ij} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_{iI} \frac{Z_I}{|\mathbf{r}_i - \mathbf{R}_I|}$$



Electronic Schrödinger equation (second quantization)



- Occupation numbers $\mathbf{n}, n_k = 0, 1, \sum_k n_k = N, \mathbf{n} \equiv \{k_1, \dots, k_N\}$
- Antisymmetric basis states are *Slater determinants*, $|\mathbf{n}\rangle = \det \varphi_{k_i}(\mathbf{r}_i)$, (asymmetric) wave function $\psi(\mathbf{n}) \equiv c_{\mathbf{n}}$
- Local energy as a sum over all Slater determinants that differ in occupation of up to two orbitals

$$E_{\text{loc}}(\mathbf{n}) = \sum_{\mathbf{n}':k\to l} h_{kl} \frac{\psi(\mathbf{n}')}{\psi(\mathbf{n})} + \sum_{\mathbf{n}'':k,l\to m,n} V_{klmn} \frac{\psi(\mathbf{n}'')}{\psi(\mathbf{n})}$$
$$h_{kl} = \int d\mathbf{r} \,\varphi_k(\mathbf{r}) \left(-\frac{1}{2}\Delta - \sum_I \frac{Z_I}{|\mathbf{r} - \mathbf{R}_I|} \right) \varphi_l(\mathbf{r})$$
$$V_{klmn} = \frac{1}{2} \int d\mathbf{r} d\mathbf{r}' \frac{\varphi_k(\mathbf{r})\varphi_l(\mathbf{r}')\varphi_m(\mathbf{r})\varphi_n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}$$

Determinant expansion curse



Example 1: PauliNet

nature chemistry



Deep-neural-network solution of the electronic Schrödinger equation

Jan Hermann^{1,2}, Zeno Schätzle¹ and Frank Noé^{1,3,4}

1. Slater determinant

 $\psi_{\theta}(\mathbf{r}) := \det \varphi_k(\mathbf{r}_i), \qquad k = 1, \dots, N$

2. Generalized Slater determinant

 $\psi_{\theta}(\mathbf{r}) := \det f_{ki,\theta}(\mathbf{r}), \qquad \mathbf{f}_{\theta}(\mathcal{P}_{ij}\mathbf{r}) = \mathcal{P}_{ij}\mathbf{f}_{\theta}(\mathbf{r})$

3. Real-space baseline/envelope

 $\psi_{\boldsymbol{\theta}}(\mathbf{r}) := \det f_{ki,\boldsymbol{\theta}}(\mathbf{r}) \boldsymbol{\varphi}_k(\mathbf{r}_i)$





ARTICLE

Example 2:

quantum states

Neural

https://doi.org/10.1038/s41467-020-15724-9



Fermionic neural-network states for ab-initio electronic structure

Kenny Choo^{1 \boxtimes}, Antonio Mezzacapo^{2 \boxtimes} & Giuseppe Carleo^{3 \boxtimes}

• Instead of storing a CI vector $c_{\mathbf{n}}$ (FCI), sampling it stochastically (FCIQMC), compressing it (CC), represent it with a NN, $\psi_{\theta}(\mathbf{n})$







Summary

- Machine learning can help avoiding as well as solving electronic structure
- Two leading ML techniques: Neural networks and Gaussian processes
- Machine-learning force fields are a mature tool ready for production calculations
- Machine-learned Hamiltonians and wave functions are a promising tool, but still in the making

SchNet for molecules

- Instance of a graph neural network
- Original SchNet is a force field $E(\{Z_i, \mathbf{R}_i\})$

$$\mathbf{x}_{i}^{(0)} := \mathbf{X}_{\boldsymbol{\theta}, Z_{i}}$$

$$\mathbf{z}_{i}^{(n)} := \sum_{j} \mathbf{w}_{\boldsymbol{\theta}}^{(n)} \left(\mathbf{e}(|\mathbf{R}_{i} - \mathbf{R}_{j}|) \right) \odot \mathbf{h}_{\boldsymbol{\theta}}^{(n)} \left(\mathbf{x}_{j}^{(n)} \right)$$

$$\mathbf{x}_{i}^{(n+1)} := \mathbf{x}_{i}^{(n)} + \mathbf{g}_{\boldsymbol{\theta}}^{(n)} \left(\mathbf{z}_{i}^{(n)} \right)$$

$$E := \sum_{i} \varepsilon_{\boldsymbol{\theta}} \left(\mathbf{x}_{i}^{(N)} \right)$$

Schütt et al. J. Chem. Phys. 148, 241722 (2018)

Graph neural networks initialize $\mathbf{x}_{i}^{(0)}$ $\mathbf{z}_{i}^{(n)} \coloneqq \sum_{j} \mathbf{p}_{\theta}^{(n)}(\mathbf{x}_{i}^{(n)}, \mathbf{x}_{j}^{(n)}, \mathbf{e}_{ij})$ $\mathbf{x}_{i}^{(n+1)} \coloneqq \mathbf{q}_{\theta}^{(n)}(\mathbf{x}_{i}^{(n)}, \mathbf{z}_{i}^{(n)})$ $\mathcal{O}(N_{edge}^{2})$



SchNet for electrons in molecules

- Nodes represent electrons
- Separate same-spin, opposite-spin, and nuclear messages
- Single SchNet instance for both the Jastrow factor $J_{\theta}(\mathbf{r})$ and backflow $\mathbf{f}_{\theta}(\mathbf{r})$

$$\begin{aligned} \mathbf{x}_{i}^{(0)} &:= \mathbf{X}_{\boldsymbol{\theta}, s_{i}} \\ \mathbf{z}_{i}^{(n, \pm)} &:= \sum_{j \neq i}^{\pm} \mathbf{w}_{\boldsymbol{\theta}}^{(n, \pm)} \left(\mathbf{e}(|\mathbf{r}_{i} - \mathbf{r}_{j}|) \right) \odot \mathbf{h}_{\boldsymbol{\theta}}^{(n)} \left(\mathbf{x}_{j}^{(n)} \right) \\ \mathbf{z}_{i}^{(n, n)} &:= \sum_{J} \mathbf{w}_{\boldsymbol{\theta}}^{(n, n)} \left(\mathbf{e}(|\mathbf{r}_{i} - \mathbf{R}_{J}|) \right) \odot \mathbf{Y}_{\boldsymbol{\theta}, J} \\ \mathbf{x}_{i}^{(n+1)} &:= \mathbf{x}_{i}^{(n)} + \sum_{\pm} \mathbf{g}_{\boldsymbol{\theta}}^{(n, \pm)} \left(\mathbf{z}_{i}^{(n, \pm)} \right) + \mathbf{g}_{\boldsymbol{\theta}}^{(n, n)} \left(\mathbf{z}_{i}^{(n, n)} \right) \\ J &:= \eta_{\boldsymbol{\theta}} \left(\sum_{i} \mathbf{x}_{i}^{(N)} \right) \\ \mathbf{f}_{i} &:= \kappa_{\boldsymbol{\theta}} \left(\mathbf{x}_{i}^{(N)} \right) \end{aligned}$$
 Graph neural network