



How Machine Learning and Chemistry Meet in the Quantum Chemistry Field.

Johnatan Mucelini

Advisor: Juarez L. F. Da Silva

August, 2020

Quantum Chemistry

Obtaining Properties with Quantum Chemistry:

System: $\{\mathbf{R}_i, Z_i\}, N$

\mathbf{R}_i : nuclear positions

Z_i : nuclear charge

N : # electrons



Mapping into energy: $\{\mathbf{R}_i, Z_i\} \xrightarrow{\Psi} E$

QC equations:

WF: Schrödinger equation $\hat{H}\Psi = E\Psi$

DFT: Kohn-Sham equation $\hat{H}_{\text{eff}}\phi_{ks}^i = e^i\phi_{ks}^i$
 $\rho = \sum_i |\phi_{ks}^i|^2$

Mapping into other properties: $E, \rho \rightarrow \text{property}$

$$F = \frac{\partial E}{\partial R_i} \quad Q_i = \int_{\text{basis } i} \rho d\tau - Z_i$$

Problems related to solving QC equation:

- Exponential scaling of computational cost for large systems:
 DFT (cheap method): approximately $O(N_e^3)$
 DFT (VASP): close to $O(N_{atoms}^2 \ln N_{atoms})$
 Computational Materials Science 6 (1996) 15-50
- In global optimizations and high-throughput screening:
 - QC equations are extensively solved.
 - Big data complicate some tasks, such as extract representative sets.

Dream: functions that map systems ($\{\mathbf{R}_i, Z_i\}$) into the energy and properties directly: (without QC equations)

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

Representation, how ML “see” QC

ML “see” QC system (take the input) as a vector of values (features) that represent the system.

$$\mathbf{x} = [x_1, x_2, \dots]$$

Intuitive descriptors:

- $\{\mathbf{R}_i, Z_i\}$.
- # bonds, average/ranked bond distance.

Qualities of a good representation:

- Invariance to the rotation, translation, and homo-nuclear permutational.
- Non-degenerate.
- Unique.
- Size extensively.

Smooth Overlap of Atomic Positions (SOAP)
Neighbor density function:

$$\rho^j = \sum_j \exp(-\sigma|\mathbf{r} - \mathbf{r}_j|^2) = \sum_{nlm} c_{nlm}^i g_n(r) Y_{lm}(\hat{\mathbf{r}})$$

where $g_n(r)$ is a orthonormal radial basis function.

$$\mathbf{X} = \mathbf{Y} = \mathbf{Y} \quad \mathbf{Y} = \mathbf{X} = \mathbf{Y}$$

$$\rho^{\mathbf{X}} \quad \mathbf{X} \quad \text{⊗} \quad \text{⊗} \quad \text{⊗} \quad \mathbf{X} \quad \text{⊗}$$

From $\rho^{\mathbf{X}}$ one can derive the rotational invariants

$$x_{nn'l} = \sum_m c_{nlm} (c_{n'l m})^*$$

which can be employed as elements of a descriptor vector $\mathbf{x}^{\mathbf{X}}$, and $\mathbf{x} = [\mathbf{x}^{\mathbf{X}}, \mathbf{x}^{\mathbf{Y}}, \mathbf{x}^{\mathbf{Y}'}]$.

Paradigms, how ML learn from QC examples/data

Two main classes: **Supervised** and **Unsupervised learning**

Supervised learning, simple regression:

Input: \mathbf{X} (features), \mathbf{y} (target/label)

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}^1 \\ \mathbf{x}^2 \\ \dots \\ \mathbf{x}^N \end{bmatrix} = \begin{bmatrix} x_1^1 & x_2^1 & \dots & x_M^1 \\ x_1^2 & x_2^2 & \dots & x_M^2 \\ \dots & \dots & \dots & \dots \\ x_1^N & x_2^N & \dots & x_M^N \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} y^1 \\ y^2 \\ \dots \\ y^N \end{bmatrix}$$

$h_{\theta}(\mathbf{x}^i)$: model (map $\mathbf{x}^i \rightarrow y^i$)

$J(\mathbf{y}, h_{\theta}(\mathbf{X}))$: cost function


$$J = \sum_i^N |y^i - h_{\theta}(\mathbf{x}^i)|^2$$

Learning task: $\operatorname{argmin}_{\theta} \{l_n(\mathbf{y}, h_{\theta}(\mathbf{X}))\}$

Input: $\{y^i, x^i\}$ (\mathbf{x}^i present 1 element/feature)

Polynomial model: $h_{\theta}(x) = \theta_0 + \theta_1 x^1 + \theta_2 x^2 + \theta_3 x^3$

Cost function: $J = \sum_i |y^i - h_{\theta}(x^i)|^2$



Training set: $X^{train} = \mathbf{X}[70\%]$

Test set: $X^{test} = \mathbf{X}[30\%]$

Paradigms, how ML learn from QC examples/data

Two main classes: **Supervised** and **Unsupervised learning**

Unsupervised learning, KMeans clusterization:

Input: \mathbf{X} (features)


$$\mathbf{X} = \begin{bmatrix} \mathbf{x}^1 \\ \mathbf{x}^2 \\ \dots \\ \mathbf{x}^N \end{bmatrix} = \begin{bmatrix} x_1^1 & x_2^1 & \dots & x_M^1 \\ x_1^2 & x_2^2 & \dots & x_M^2 \\ \dots & \dots & \dots & \dots \\ x_1^N & x_2^N & \dots & x_M^N \end{bmatrix}$$

For a given number of $\mathbf{C}_i \in \mathbb{R}^M$

Algorithm:

- 0) Initialize \mathbf{C}_i .
- 1) \mathbf{x}^i is associated with the closest \mathbf{C}_j .
- 2) \mathbf{C}_j is the mean point of the \mathbf{x}^i associated.
- 3) Repeat steps 1 and 2 until convergence.

Problems: Fall in local minimums, equal clusters sizes, sphere shape




Machine Learning: Neural Networks


The Perceptron:

$$f(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{w} \cdot \mathbf{x} + b \geq 0 \\ 0 & \text{else} \end{cases}$$

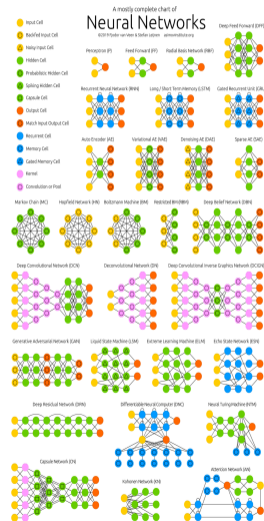

The Neural Network node:



Activation functions:



A Neural Network:



Gaussian Process Regression


Bayesian Linear Regression

$$h(\mathbf{x}) = \mathbf{x}^T \mathbf{w}, \quad y = h(\mathbf{x}) + \varepsilon, \quad \varepsilon \sim N(0, \sigma_n^2)$$

$$p(\mathbf{w} | \mathbf{y}, \mathbf{X}) = \frac{p(\mathbf{y} | \mathbf{X}, \mathbf{w}) p(\mathbf{w})}{p(\mathbf{y} | \mathbf{X})}, \quad p(\mathbf{w}) \sim N(0, \Sigma)$$

where $A = \sigma^{-2} \mathbf{X} \mathbf{X}^T + \Sigma^{-1}$. Predicting h for \mathbf{x}' :

$$p(h(\mathbf{x}') | \mathbf{x}', \mathbf{X}, \mathbf{y}) = N(\mathbf{x}'^T A^{-1} \mathbf{X} \mathbf{y}, \mathbf{x}'^T A^{-1} \mathbf{x}')$$



Seppe
vanden
Broucke et
al.,
medium.com,
2018

The Kernel Trick:

$\phi(\mathbf{x})$ maps \mathbf{x} into another space. Defining $\psi(\mathbf{x}) = \Sigma^{1/2} \phi(\mathbf{x})$ we obtain a dot product $\psi(\mathbf{x}^i) \cdot \psi(\mathbf{x}^j) = \phi(\mathbf{x}^i)^T \Sigma \phi(\mathbf{x}^j) = k(\mathbf{x}^i, \mathbf{x}^j) = k_{ij} \in \mathfrak{R}$.

Gaussian Process Regression

Applying Φ on \mathbf{x} , $h(\mathbf{x}) = \phi(\mathbf{x})^T \mathbf{w}$, one found the result with k instead ϕ :


$$p(h | \mathbf{x}', \mathbf{X}, \mathbf{y}) = N(\phi'^T \Sigma \Phi (K + \sigma^2 I)^{-1} \mathbf{y}, S)$$

$$S = \phi'^T \Sigma \phi' - \phi'^T \Sigma \Phi (K + \Sigma^2 I)^{-1} \Phi^T \sigma \phi'$$


where $K = \Phi^T \Sigma \Phi$ and $\Phi = \Phi(\mathbf{X})$.

$$K = \begin{bmatrix} k_{0,0} & k_{0,1} & \dots & k_{0,N} \\ k_{1,0} & k_{1,1} & \dots & k_{1,N} \\ \dots & \dots & \dots & \dots \\ k_{N,0} & k_{N,1} & \dots & k_{N,N} \end{bmatrix}$$

Gaussian Process Regression



(a), prior



(b), posterior





C. E.
Rasmussen &
C. K. I.
Williams,
Gaussian
Processes for
Machine
Learning,
2006.

Examples: Kolsbjerg et al., 2018

ML + Evolutionary algorithm + QC:

- ML Model: NN (2 hidden layers, 5 nodes each).
- Representation: Gaussian (G2 and G4) descriptor.
- Finding the global minimum structure with 50% certainty per run:
 - no NN: ~ 170 generations.
 - +NN: ~ 260 generations.
 - +NN runs successful is delayed.
- Successful runs:
 - no NN: ~ 8900 QC calculations.
 - +NN: ~ 260 QC calculations.
 - +NN reduce # calculations by $\times 40$.

Phys. Rev. B, 97, 2018, 195424



Examples: Podryabinkin et al., 2019

ML + Evolutionary algorithm + QC:

- ML Model: Moment Tensor Potentials (MTP).

$$E = \sum_i V_i$$

$$V_i = \sum_j w_j B_j(\mathbf{u}_i)$$

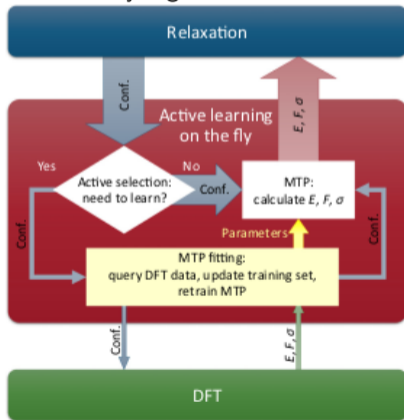
- Evolutionary algorithm: USPEX.

- Active learning: prediction “extrapolation”.


- Time spend to find the 6 lowest-energy structures reduce 1-4 orders.

- Efficiency is improved with pre-training.


- Configuration relaxation inside the evolutionary algorithm:



- Material: Boron allotropes




β -boron approximant
 $E^{\text{DFT}} = -6.704$ eV/atom,
 Atoms: 106,
 Space group: $P1$,
 Core-hours: $7 \cdot 10^3$ AL-MTP vs. $6.6 \cdot 10^7$ DFT
 $|E^{\text{DFT}} - E^{\text{MTP}}| = 10.1$ meV/atom




Phys. Rev. B 99, 064114 (2019)

Examples: Gómez-Bombarelli et al., 2018



- Variational Autoencoder (VAE): a NN used for dimensionality reduction and generative processes.
- SMILES: string representation of molecules.
- Predictor model: Gaussian Process
- Autoencoder training: 250k molecules (ZINC)
- Encodes approximately: 7.5M molecules



ACS Cent. Sci. 4, 2018, 268-276

BRIEF COMMUNICATION

<https://doi.org/10.1038/s41587-019-0224-x>

nature
biotechnology

Deep learning enables rapid identification of potent DDR1 kinase inhibitors

Alex Zhavoronkov^{1*}, Yan A. Ivanenkov¹, Alex Aliper¹, Mark S. Veselov¹, Vladimir A. Aladinskiy¹, Anastasiya V. Aladinskaya¹, Victor A. Terentiev¹, Daniil A. Polykovskiy¹, Maksim D. Kuznetsov¹, Arip Asadulaev¹, Yury Volkov¹, Artem Zholus¹, Rim R. Shayakhmetov¹, Alexander Zhebrak¹, Lidiya I. Minaeva¹, Bogdan A. Zagribelnyy¹, Lennart H. Lee², Richard Soll², David Madge², Li Xing², Tao Guo² and Alán Aspuru-Guzik^{3,4,5,6}

Nat. Biotechnol., 37, 2019, 1038-1040

Overview and Conclusions

- Quantum Chemistry.
 - High computational costs.
 - Many calculations are necessary.
- Machine Learning.
 - Representation.
 - How ML “see” chemistry structures.
 - Supervised:
 - Ex.: NN, gaussian process.
 - Unsupervised:
 - Ex.: clustering, VAE.
- Wide range of applications:
 - Ex.: Nanomaterials and Drug Design.
 - Commonly coupled with other algorithms.

Acknowledgments



Thanks for your Attention!

Supporting Slides