
HydroNet: Benchmark Tasks for Preserving Intermolecular Interactions and Structural Motifs in Predictive and Generative Models for Molecular Data

Sutanay Choudhury
sutanay.choudhury@pnnl.gov

Jenna A. Bilbrey
jenna.bilbrey@pnnl.gov

Logan Ward
lward@anl.gov

Sotiris S. Xantheas
sotiris.xantheas@pnnl.gov

Ian Foster
foster@anl.gov

Joseph P. Heindel
heindelj@uw.edu

Ben Blaiszik
blaiszik@uchicago.edu

Marcus E. Schwarting
meschw04@uchicago.edu

Abstract

Intermolecular and long-range interactions are central to phenomena as diverse as gene regulation, topological states of quantum materials, electrolyte transport in batteries, and the universal solvation properties of water. We present a set of challenge problems¹ for preserving intermolecular interactions and structural motifs in machine-learning approaches to chemical problems, through the use of a recently published dataset of 4.95 million water clusters held together by hydrogen bonding interactions and resulting in longer range structural patterns. The dataset provides spatial coordinates as well as two types of graph representations, to accommodate a variety of machine-learning practices.

1 Introduction

The application of machine-learning (ML) techniques such as supervised learning and generative models in chemistry is an active research area. ML-driven prediction of chemical properties and generation of molecular structures with tailored properties have emerged as attractive alternatives to expensive computational methods [20, 24, 23, 32, 31, 7, 14, 16, 22]. Though increasingly used, graph representations of molecules often do not explicitly include non-covalent interactions such as hydrogen bonding, which poses difficulties when examining systems with intermolecular and/or long-range interactions [10]. To facilitate the development of such methods, we discuss a set of challenge problems and suggest an approach based on a recently published database of low-energy water cluster minima lying within 5 kcal/mol from the putative minimum of each cluster size [17].

Scientific motivation: A water cluster is a discrete hydrogen bonded network of water molecules. While most interactions are short-range (i.e., between neighboring molecules) [19], there also exist substantial (~20%) many-body, longer-range interactions (i.e., with next-nearest and more distant neighbors) [30]. Understanding many-body and long-range hydrogen bonding interactions is key to answering long-standing scientific questions such as the macroscopic properties of liquid water, ice, and aqueous systems (e.g., heat capacity, density, dielectric constant, compressibility) [13]. These interactions are responsible for the bulk and interfacial properties of liquid water and ice, as well as solvation processes, and are key to the realization of diverse applications, from drug delivery to protein folding and the design of quantum materials and novel electrolytes for batteries [15, 28, 9].

¹<https://exalearn.github.io/hydronet>

HydroNet summary: The dataset of 4.95 million water cluster minima is the largest collection of water cluster minima reported to date [17]. Originally created to advance the development of interaction potentials in chemical physics, the dataset is composed of clusters of isomers differing in the underlying hydrogen bonding network. Each cluster is described by its potential energy, Cartesian coordinates of each atom, and two graph-based representations of their bonding arrangement (Fig. 1): an *atomic interaction graph* that captures both intramolecular and intermolecular bonding patterns and a *coarse graph* that captures only the intermolecular structure of the cluster.




Figure 1: Data model: geometries provided as spatial coordinates. In the graph representation nodes are atoms (Atomic) or water molecules (Coarse) and edges are covalent and/or hydrogen bonds.

ML tasks on the dataset: 1) *Molecular Property Prediction:* given a water cluster with specified spatial coordinate information or bonding structure, predict its energy. 2) *Molecule Generation:* given N water molecules, generate candidate structures that conform to structural measures of low-energy configurations described in the key challenges below.

Key challenges: A defining feature of water clusters is that numerous dissimilar structures can have quite similar energies. In addition, for a given set of spatial orientations of oxygen atoms in the cluster (oxygen network), there exist numerous hydrogen bonding networks (Coarse graphs) depending on the arrangement of the hydrogen atoms, which form hydrogen bonds according to the Bernal–Fowler rules [3]. The structural properties of low-energy hydrogen bonding networks – characterized by graph-theoretical measures such as degree distribution, shortest path length, and distribution of polygons (Fig. 2) – vary systematically with cluster size [4]. A generative method that produces a water cluster network should be mindful of these properties, as clusters far outside the distributions are likely to be much higher in energy and therefore of less interest.




Figure 2: Graph-based measures capturing intermolecular interactions and structural motifs in water cluster networks. Clear patterns such as the average number of neighbors (left), the shortest path between two water molecules (middle), and the predominance of pentamer and hexamer polygons over tetramer polygons (right) emerge as the cluster size increases [4].

2 Dataset Description

The potential energy of each cluster was obtained using the ab initio-based Thole-type, flexible, polarizable interaction potential for water (TTM2.1-F) [6, 8]. Clusters were generated with the Monte Carlo Temperature Basin Paving (MCTBP) sampling method to produce a dense sampling of low-energy water clusters containing 3–30 water molecules per cluster [17]. The water cluster minima dataset [2] is represented in the three formats shown in Fig. 1, each as both line-delimited JavaScript Object Notation (JSON) and Tensorflow Protobufs. Table 1 lists the information available for each sample. The records are separated into predefined train (80%), validation (10%), and test sets (10%), where clusters are maintained in the same subset for each data format. We also provide code to compute graph descriptors for structural motif tracking [1].

3 Machine-Learning Tasks

We introduce the defined property prediction and generative modeling tasks on the water cluster dataset and provide details on baseline implementations for the first task.

Table 1: Records in HydroNet are stored as dictionary objects with the below keys.

Key	Type	Description
energy	float	Energy of the cluster, in kcal/mol.
n_water	int	Number of water molecules.
n_atom	int	Number of atoms or, for coarse graphs, nodes.
atom	list of ints	Types of each atom or node in the graph. Atomic: 0 is oxygen, 1 is hydrogen; coarse: all nodes are the same type.
z	list of ints	<i>Geometry only.</i> Atomic number of each atom.
coords	$N \times 3$ array of floats	<i>Geometry only.</i> XYZ coordinates of atoms, in angstrom.
n_bond	list of ints	<i>Graph only.</i> Number of edges in the graph.
bond	list of ints	<i>Graph only.</i> Types of edges. Atomic: 0 is covalent, 1 is hydrogen; coarse: 0 is donor, 1 is acceptor.
connectivity	$N \times 2$ array of ints	<i>Graph only.</i> Connectivity between nodes in the graph.

3.1 Cluster Potential Energy Prediction Task

This task is to predict the potential energy of a water cluster without the use of expensive ab initio methods, in two settings. In the *geometry-to-energy* [21] setting, we assume the geometry of the structure is known accurately. In the *graph-to-energy* [11, 26] setting, we require predictions to be made from the connectivity of the water molecules only.

We approach these tasks using neural network potentials, in which the original state of each atom (h_v) and bond (α_{vw}) is represented as a vector embedding based on atomic number and bond type. These states are modified by successive message layers. Each message layer uses a multi-layer perceptron to compute a message from the atom state (h_v), a neighboring atom state (h_w), and the connecting bond (α_{vw}). The atom and bond states are updated according to the following equations, which generalize both the geometry-to-energy and graph-to-energy prediction settings described below. In the first, an atom’s neighbors are those within a radial cutoff distance; in the second, neighbors are represented by the adjacency matrix of the network.

$$m_v^{t+1} = \sum_{w \in \text{Nbors}(v)} M_t(h_v^t, h_w^t, \alpha_{vw}^t) \quad h_v^{t+1} = h_v^t + m_v^{t+1} \quad \alpha_{vw}^{t+1} = \alpha_{vw}^t + M_t(h_v^t, h_w^t, \alpha_{vw}^t)$$

We sum the energy per atom values produced by the penultimate layer to obtain the full cluster energy.

For the geometry-to-energy problem, we previously [4] trained a continuous-filter convolutional neural network (SchNet, PyTorch implementation [20]) on 500,000 water clusters of size $N=11-29$. The training set was stratified by cluster size, from which 10% were reserved for validation during training. The test set was composed of 10,500 clusters, with 500 from each cluster size $N=10-30$. Clusters of size $N=10, 30$ allowed us to examine the effects of energy prediction on clusters smaller and larger than those in the training set. Using mean squared error (MSE) as the loss, we achieved a final training loss of $0.0030 \text{ (kcal/mol)}^2$ and final validation loss of $0.0035 \text{ (kcal/mol)}^2$. The mean absolute error (MAE) of the test set predictions was 0.0427 kcal/mol . The network took ~ 65 hours to train using 4 NVIDIA V100 GPUs.

For the graph-to-energy problem, we use a message-passing neural network (MPNN) in-

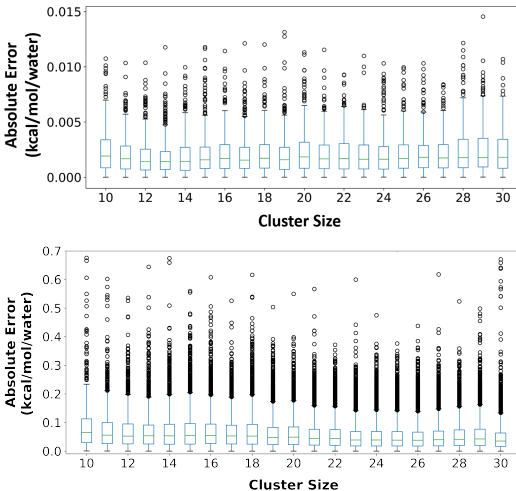


Figure 3: Performance of (top) SchNet trained on 500,000 water clusters, $N=11-29$, and tested on 10,500 clusters, $N=10-30$, and (bottom) MPNN trained on HydroNet training set and evaluated on test set. SchNet performance data from [4].

spired by [11] to create a reference implementation [26]. A first performance study shows that models trained on coarse networks achieve losses an order of magnitude superior to those trained on atomic networks, but exhibit significantly worse performance on the test set. Such generalization issues suggest opportunities to explore regularization techniques or new network designs to take advantage of coarse graphs.

Our best models have a MAE of ~ 0.5 kcal/mol/water. Errors are nearly equal for clusters of sizes 10–30, giving us confidence that the MPNN is capturing the changes in favored geometries as clusters increase in size. However, these errors are around a factor of 100 higher than those achieved with SchNet, which incorporates knowledge of the atomic coordinates. In short, we find conventional MPNN approaches to be promising solutions for energy prediction, but significant work is needed by the ML community to find models suitable for evaluating systems with intermolecular and/or long-range interactions.

3.2 Structural Measure-Preserving Molecular Generation Task

We define the task as follows: given N water molecules, generate a geometric/atomic/coarse representation that a) satisfies certain graph-theoretic structural measures derived from its hydrogen bonding network and b) minimizes the cluster energy by optimizing the relative spatial arrangements of atoms and molecules. As an example, Fig. 4 shows a sequence of transformations that would yield a more stable water cluster from a static number of molecules.

Structural measures vary with cluster size, as shown in Fig. 2. The mean degree (number of neighbors) of water molecules in a hydrogen bonded network increases and eventually plateaus when water molecules become saturated with hydrogen bonds. The shortest path length also increases with cluster size, through the growth rate decays. The number of cycles in the hydrogen bonded network also show trends as clusters grow in size: trimers (3-cycles) and tetramers (4-cycles) become less dominant, while pentamers (5-cycles) and hexamers (6-cycles) emerge as key building blocks. We also observe a strong evolution in geometric structure (Fig. 5). The water cluster structures start as quasi-planar graphs ($N=3-5$), evolve into cubic graphs ($N=7-16$), and then into cage-like structures ($N=17$ onward), with some structures being quasi-symmetric (i.e. for $N=20$).




Figure 4: Generation of a low-energy water network from a higher energy structure.

4 Relationship with Other Work

Generative models and deep reinforcement learning methods that produce molecular structures with desired properties are active research areas [31, 32, 7, 12, 16, 22]. However, most ML-benchmark datasets [18, 27, 29] used by these studies map from spatial or graph structure to properties of interest (e.g. drug-likeness, binding affinity between a drug molecule and a target protein). Modeling the interaction of heterogeneous functional groups and reasoning about potential reaction mechanisms is the primary focus for such tasks. Our dataset and science challenge tasks require methods that prioritize satisfying topological constraints, addressing challenges with stereoisomers, and modeling variations in interaction patterns at multiple scales.

The multi-scale nature of the interactions of our dataset also fills a void in existing scientific datasets. Datasets which map molecular geometry or molecular graphs, available for years, capture challenges in which properties are driven by short-range covalent interactions [25, 29, 18]. There are also datasets of condensed phases (e.g., AGNI [5]) in which long-range, many-body interactions become important. Our dataset presents a complex mixture of short-range covalent, intermolecular hydrogen bonding, and extended many-body effects. We propose this dataset as a basis for the challenges of learning such effects from 3D geometries and inferring them from bonding structure.




Figure 5: Geometric shape progression of low-energy structures with cluster size n (figure reproduced from [17]).

5 Conclusions

We present challenge problems and a corresponding dataset to advance machine-learning on molecular data. An important feature of the dataset discussed here is the incorporation of intermolecular interactions, which allows for two distinct graph representations in addition to the spatial coordinates. Our initial benchmarks illustrate a large performance gap between neural networks based on atomic coordinates and those which use only the bonding structure, which make this benchmark particularly relevant for advancing graph neural network-based generative and prediction models. Moreover, this dataset is useful for the creation of novel machine-learning methods that preserve structural patterns across a wide size regime and can address the key challenges of property prediction and molecular generation.

6 Broader Impact

The problems addressed here have the potential to advance methods for modeling biological and chemical systems, leading to vast speed-ups in computation that would accelerate accurate simulations of large, complex systems. These speed-ups will allow domain researchers to more quickly assess the properties of and generate potential structures for novel molecular systems with phenomena dependant on long-range interactions. Such research is the beginning stage to making practicable, real-world scientific advancements.

7 Acknowledgements

S.C., J.A.B., L.W., S.S.X., I.F., J.P.H., and M.S. were supported by the DOE Exascale Computing Project, ExaLearn Co-design Center. B.B. was supported by the National Science Foundation under NSF Award Number: 1931306 “Collaborative Research: Framework: Machine Learning Materials Innovation Infrastructure.” This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. Pacific Northwest National Laboratory (PNNL) is a multi-program national laboratory operated for the US Department of Energy by Battelle. Argonne National Laboratory’s work was supported by the U.S. Department of Energy, Assistant Secretary for Environmental Management, Office of Science and Technology, under contract DE-AC02-06CH11357.

References

- [1] Code repository for computing structural motifs. <https://github.com/exalearn/molecular-graph-descriptors>.
- [2] Database of water clusters. <https://sites.uw.edu/wdbase>.
- [3] J. D. Bernal and R. H. Fowler. A theory of water and ionic solution, with particular reference to hydrogen and hydroxyl ions. *The Journal of Chemical Physics*, 1(8):515–548, 1933.
- [4] Jenna A. Bilbrey, Joseph P. Heindel, Malachi Schram, Pradipta Bandyopadhyay, Sotiris S. Xantheas, and Sutanay Choudhury. A look inside the black box: Using graph-theoretical descriptors to interpret a Continuous-Filter Convolutional Neural Network (CF-CNN) trained on the global and local minimum energy structures of neutral water clusters. *The Journal of Chemical Physics*, 153(2):024302, 2020.
- [5] V. Botu, R. Batra, J. Chapman, and R. Ramprasad. Machine learning force fields: Construction, validation, and outlook. *The Journal of Physical Chemistry C*, 121(1):511–522, December 2016.
- [6] Christian J Burnham and Sotiris S Xantheas. Development of transferable interaction models for water. IV. A flexible, all-atom polarizable potential (TTM2-F) based on geometry dependent charges derived from an ab initio monomer dipole moment surface. *The Journal of Chemical Physics*, 116(12):5115–5124, 2002.
- [7] Nicola De Cao and Thomas Kipf. Molgan: An implicit generative model for small molecular graphs. *arXiv preprint arXiv:1805.11973*, 2018.
- [8] George S Fanourgakis and Sotiris S Xantheas. The flexible, polarizable, Thole-type interaction potential for water (TTM2-F) revisited. *The Journal of Physical Chemistry A*, 110(11):4100–4106, 2006.
- [9] Roger H French, V Adrian Parsegian, Rudolf Podgornik, Rick F Rajter, Anand Jagota, Jian Luo, Dilip Asthagiri, Manoj K Chaudhury, Yet-ming Chiang, Steve Granick, et al. Long range interactions in nanoscale science. *Reviews of Modern Physics*, 82(2):1887, 2010.
- [10] Niklas Gebauer, Michael Gastegger, and Kristof Schütt. Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules. In *Advances in Neural Information Processing Systems*, pages 7566–7578, 2019.
- [11] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.
- [12] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Composing molecules with multiple property constraints. *arXiv preprint arXiv:2002.03244*, 2020.
- [13] Frank N Keutsch and Richard J Saykally. Water clusters: Untangling the mysteries of the liquid, one molecule at a time. *Proceedings of the National Academy of Sciences*, 98(19):10533–10540, 2001.
- [14] Yash Khemchandani, Steve O’Hagan, Soumitra Samanta, Neil Swainston, Timothy J Roberts, Danushka Bollegala, and Douglas B Kell. Deepgraphmol, a multi-objective, computational strategy for generating molecules with desirable properties: a graph convolution and reinforcement learning approach. *bioRxiv*, 2020.
- [15] Erez Lieberman-Aiden, Nynke L Van Berkum, Louise Williams, Maxim Imakaev, Tobias Ragozcy, Agnes Telling, Ido Amit, Bryan R Lajoie, Peter J Sabo, Michael O Dorschner, et al. Comprehensive mapping of long-range interactions reveals folding principles of the human genome. *Science*, 326(5950):289–293, 2009.
- [16] Kaushalya Madhawa, Katushiko Ishiguro, Kosuke Nakago, and Motoki Abe. Graphnvp: An invertible flow model for generating molecular graphs. *arXiv preprint arXiv:1905.11600*, 2019.

- [17] Avijit Rakshit, Pradipta Bandyopadhyay, Joseph P. Heindel, and Sotiris S. Xantheas. Atlas of putative minima and low-lying energy networks of water clusters $n = 3-25$. *The Journal of Chemical Physics*, 151(21):214307, 2019.
- [18] Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1, 2014.
- [19] Laurent Schaeffer. Chapter 14 - the role of functional groups in drug-receptor interactions. In Camille Georges Wermuth, David Aldous, Pierre Raboisson, and Didier Rognan, editors, *The Practice of Medicinal Chemistry (Fourth Edition)*, pages 359 – 378. Academic Press, San Diego, fourth edition edition, 2008.
- [20] K. T. Schütt, P. Kessel, M. Gastegger, K. A. Nicoli, A. Tkatchenko, and K.-R. Müller. Schnet-pack: A deep learning toolbox for atomistic systems. *Journal of Chemical Theory and Computation*, 15(1):448–455, 2019.
- [21] K. T. Schütt, H. E. Saucedo, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller. Schnet – a deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24):241722, 2018.
- [22] Chence Shi, Minkai Xu, Zhaocheng Zhu, Weinan Zhang, Ming Zhang, and Jian Tang. Graphaf: a flow-based autoregressive model for molecular graph generation. *ICLR*, 2020.
- [23] Andreas Singraber, Tobias Morawietz, Jörg Behler, and Christoph Dellago. Parallel multistream training of high-dimensional neural network potentials. *Journal of Chemical Theory and Computation*, 15(5):3075–3092, 2019. PMID: 30995035.
- [24] J. S. Smith, O. Isayev, and A. E. Roitberg. Ani-1: an extensible neural network potential with dft accuracy at force field computational cost. *Chem. Sci.*, 8:3192–3203, 2017.
- [25] Justin S. Smith, Olexandr Isayev, and Adrian E. Roitberg. ANI-1, a data set of 20 million calculated off-equilibrium conformations for organic molecules. *Scientific Data*, 4(1), December 2017.
- [26] Peter C St. John, Caleb Phillips, Travis W Kemper, A Nolan Wilson, Yanfei Guan, Michael F Crowley, Mark R Nimlos, and Ross E Larsen. Message-passing neural networks for high-throughput polymer screening. *The Journal of Chemical Physics*, 150(23):234111, 2019.
- [27] Teague Sterling and John J Irwin. ZINC 15–ligand discovery for everyone. *Journal of Chemical Information and Modeling*, 55(11):2324–2337, 2015.
- [28] Nevena Veljkovic, Sanja Glisic, Vladimir Perovic, and Veljko Veljkovic. The role of long-range intermolecular interactions in discovery of new drugs. *Expert Opinion on Drug Discovery*, 6(12):1263–1270, 2011.
- [29] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. MoleculeNet: A benchmark for molecular machine learning. *Chemical Science*, 9(2):513–530, 2018.
- [30] Sotiris S Xantheas. Cooperativity and hydrogen bonding network in water clusters. *Chemical Physics*, 258(2):225 – 231, 2000.
- [31] Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. In *Advances in Neural Information Processing Systems*, pages 6410–6421, 2018.
- [32] Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N. Zare, and Patrick Riley. Optimization of molecules via deep reinforcement learning. *Scientific Reports*, 9(1):10752, 2019.