
3D Infomax improves GNNs for Molecular Property Prediction

Hannes Stärk^{*1} Dominique Beaini² Gabriele Corso¹ Prudencio Tossou² Christian Dallago³
Stephan Günnemann³ Pietro Liò⁴

Abstract

Molecular property prediction is one of the fastest-growing applications of deep learning with critical real-world impacts. However, although the 3D molecular graph structure is necessary for models to achieve strong performance on many tasks, it is infeasible to obtain 3D structures at the scale required by many real-world applications. To tackle this issue, we propose to use existing 3D molecular datasets to pre-train a model to reason about the geometry of molecules given only their 2D molecular graphs. Our method, called 3D Infomax, maximizes the mutual information between learned 3D summary vectors and the representations of a graph neural network (GNN). During fine-tuning on molecules with unknown geometry, the GNN is still able to produce implicit 3D information and uses it for downstream tasks. We show that 3D Infomax provides significant improvements for a wide range of properties, including a 22% average MAE reduction on QM9 quantum mechanical properties. Moreover, the learned representations can be effectively transferred between datasets in different molecular spaces.

1. Introduction

The understanding of molecular and quantum chemistry is a rapidly growing area for deep learning, with models having direct real-world impacts in quantum chemistry (Dral, 2020), protein structure prediction (Jumper et al., 2021), materials science (Schmidt et al., 2019), and drug discovery (Stokes et al., 2020). In particular, for the task of molecular property prediction, GNNs have had great success (Yang et al., 2019).

^{*}Equal contribution ¹EECS, Massachusetts Institute of Technology, Cambridge MA, USA ²Valence Discovery, Montreal, CA ³Department of Informatics, Technical University of Munich, DE ⁴Department of Computer Science and Technology, University of Cambridge, UK. Correspondence to: Hannes Stärk <hstark@mit.edu>.

GNNs operate on the molecular graph by updating each atom’s representation based on the atoms connected to it via covalent bonds. However, these models reason poorly about other important interatomic forces that depend on the atoms’ relative positions in space. Previous works showed that using the atoms’ 3D positions improves the accuracy of molecular property prediction (Schütt et al., 2017; Klicpera et al., 2020b; Liu et al., 2021; Klicpera et al., 2021).

However, using classical molecular dynamics simulations to explicitly compute a molecule’s geometry before predicting its properties is computationally intractable for many real-world applications. Even recent ML methods for conformation generation (Xu et al., 2021a; Shi et al., 2021; Ganea et al., 2021) are still too slow for large-scale applications. This issue, as it is summarized in Figure 1, is the motivation for our method.

Our Solution: 3D Infomax. We propose to pre-train a GNN to encode implicit 3D information in its latent vectors using publicly available molecular structures. In particular, our method, 3D Infomax, pretrains a GNN by maximizing the mutual information (MI) between its embedding of a 2D molecular graph and a learned representation of the 3D graph. This way, the GNN learns to embed latent 3D information using only the information given by the 2D molecular graphs. After pre-training, the weights can be transferred and fine-tuned on molecular tasks where no 3D information is available. For those molecules, the GNN is still able to produce implicit 3D information that can be used to inform property predictions.

Several other self-supervised learning (SSL) methods that do not use 3D information have been proposed and evaluated to pre-train GNNs and obtain better property predictions after fine-tuning (Hu et al., 2020b; You et al., 2020; Xu et al., 2021b). These often rely on augmentations (such as removing atoms) that significantly alter the molecules while assuming that their properties do not change. Meanwhile, 3D Infomax pre-training teaches the model to reason about how atoms interact in space, which is a principled and generalizable form of information.

We analyze our method’s performance by pre-training with multiple 3D datasets before fine-tuning on quantum mechanical properties and evaluating the generalization abilities of

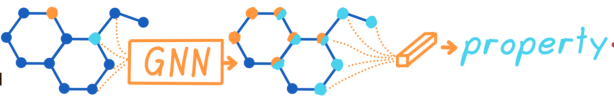



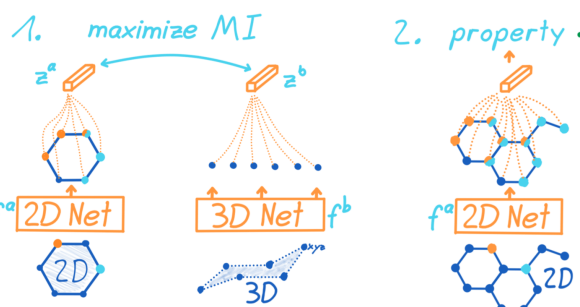


Setting	Approaches	Result
Molecules without 3D information for which properties have to be predicted.	<p>Standard Approach: Use GNNs with the molecular graph as the only input and ignore 3D based atomic interactions.</p> 	 Predictions are fast but less accurate since 3D information cannot be leveraged.
Molecules with 3D information that can be used for pre-training.	<p>Explicit 3D Approach: Employ classic (1) or learned (2) methods to compute 3D coordinates and use them as input to a 3D Graph Neural Network.</p> 	 Accurate predictions but methods for generating coordinates are too slow for many real-world applications.
	<p>Our 3D Infomax: 1. Pre-train 2D Net with the molecules for which 3D information is available and learn to generate implicit 3D information in latent representations. 2. Transfer weights of 2D Net and fine-tune for predicting molecular properties.</p> 	  During fine-tuning the 2D Net still generates latent 3D information and uses it to inform property predictions. → Predictions are more accurate than methods that do not use 3D information. → Inference is fast and only uses a single forward pass of the chosen 2D Net.

Figure 1. The different approaches to molecular property prediction and the motivation for our 3D Infomax pre-training.

the learned representations. 3D Infomax improves property predictions by large margins and the learned representations are highly generalizable: significant improvements are obtained even when the molecular space of the pre-training dataset is vastly different (e.g., in size) from the kinds of molecules in the downstream tasks. Moreover, while conventional pre-training methods sometimes suffer from negative transfer (Pan & Yang, 2010), i.e., a decrease in performance, this is not observed for 3D Infomax.

Our main contributions are:

- A 3D pre-training method that enables GNNs to reason about the geometry of molecules given only their 2D molecular graphs, which improves predictions.
- Experiments showing that our learned representations are meaningful for various molecular tasks, without negative transfer.
- Empirical evidence that the embeddings generalize across different molecular spaces.
- An approach to leverage information from multiple conformers of the same molecule that further improves downstream property predictions and an evaluation to what extent this is possible.

2. Background

2D Molecular Graphs A molecule’s 2D information can be represented as a graph $G = (\mathcal{V}, \mathcal{E})$ with atoms \mathcal{V} as nodes, and the edges \mathcal{E} given by covalent bonds. The 2D information of edges could contain the bond type while nodes are attributed with features such as the atomic number, but no 3D coordinates.

3D Molecular Conformers. Molecules are dynamic 3D structures that can exist in different spatial conformations. For a single 2D molecular graph, there are multiple low energy atom arrangements that are likely to occur in nature. These are called conformers and they can exhibit different chemical properties. For a model to properly capture 3D information, it is important to consider all the most likely conformations.

When considering a number c of known conformers of a molecule, we represent them as a set of point clouds $\{R^j\}_{j \in \{1 \dots c\}}$. Each point cloud $R = \{r_v\}_{v \in \mathcal{V}}$ specifies the locations of all atoms \mathcal{V} in the molecule.

Several tools exist to compute conformers ranging from methods based on classical force fields to slower but more accurate molecular dynamics simulations. Methods such as RDKit’s ETKDG algorithm (Landrum, 2016) are fast but less accurate. The popular metadynamics method CREST (Grimme, 2019) offers a good tradeoff between speed and ac-

curacy but still requires about 6 hours per drug-like molecule per CPU-core (Axelrod & Gomez-Bombarelli, 2020). This makes it infeasible to explicitly compute precise structures in applications where large numbers of molecules need to be processed. For instance, virtual screening for drug discovery where screening datasets sometimes are comprised of millions or billions of molecules (Gorgulla et al., 2020).

Symmetries of Molecules. A molecule’s conformation does not change if all the atom coordinates are jointly translated or rotated around a point, i.e., molecules are symmetric with respect to these two types of transformations which is also known as SE(3) symmetry. Note that some molecules (called chiral) are not invariant to reflections: their properties depend on their chirality. Deep learning architectures that capture these symmetries are usually more sample efficient, and they generalize to all symmetric inputs the architecture has been designed for (Bronstein et al., 2021). In our method, the produced representations of the 3D structure respect these symmetries.

Graph Neural Networks. We make use of GNNs to predict molecular properties given a molecular graph. Many GNNs can be described in the framework of MPNNs (Gilmer et al., 2017), such as the PNA model, (Corso et al., 2020) which we employ.

The aim of MPNNs is to learn a representation of a graph $G = (\mathcal{V}, \mathcal{E})$ with nodes \mathcal{V} connected by edges \mathcal{E} . They do so by iteratively applying message-passing layers and then combining all node representations in a readout function. A message-passing layer updates the representation of a node given its neighbors and the edges between them using permutation invariant functions such as mean, max, or sum. After the message-passing layers, another permutation invariant function can be used as a readout to obtain a final graph level representation from the node level embeddings.

3. Related Work

Molecular property prediction. Since Gilmer et al. (2017) introduced the MPNN framework, GNNs became popular for quantum chemistry (Brockschmidt, 2020; Tang et al., 2020; Withnall et al., 2020), drug discovery (Li et al., 2017; Stokes et al., 2020; Torng & Altman, 2019), and molecular property prediction in general (Coley et al., 2019; Hy et al., 2018; Unke & Meuwly, 2019). The field is well established with easily accessible molecular datasets driving progress (Wu et al., 2017; Hu et al., 2020a) and rigorous evaluations of MPNNs for property prediction (Yang et al., 2019) showing the effectiveness of the approach.

While these GNNs have had great successes by operating on the 2D graph, many tasks on molecules can be improved by additionally using 3D information. A simple approach is to use bond lengths as edge features (Chen et al., 2020a), but

methods that capture more molecular geometry improve on this such as SchNet (Schütt et al., 2017). Similarly, DimeNet (Klicpera et al., 2020b;a) proposed extracting more 3D information via bond angles, which further improved quantum mechanical property prediction. SMP (Liu et al., 2021) included another angular quantity, and GemNet (Klicpera et al., 2021) developed an approach to also capture torsion angles, such that all relative atom positions are uniquely defined. EGNN (Satorras et al., 2021) achieved the same by operating on all pairwise atom distances.

Self-Supervised Learning. (SSL) attempts to find supervision signals in unlabelled data to learn meaningful representations. Contrastive learning (van den Oord et al., 2018; Gutmann & Hyvärinen, 2010; Belghazi et al., 2018; Hjelm et al., 2019) is a popular class of methods that learn representations by comparing the embeddings of similar and dissimilar inputs and have achieved impressive results in computer vision (Chen et al., 2020b; Caron et al., 2020).

Learning from unlabeled data also is a critical challenge in molecular chemistry since datasets are relatively small due to experimental or computational costs. Several works have explored contrastive learning variants in the context of molecular graphs for non-quantum molecular properties (Hu et al., 2020b; Wang et al., 2021; You et al., 2020; 2021; Xu et al., 2021b; Zhu et al., 2021). However, the improvements these methods provide in molecular property prediction are still limited and often fail to generalize.

Previous methods for SSL on molecules only leveraged the 2D information of molecules. Meanwhile, 3D Infomax and the concurrently developed GraphMVP (Liu et al., 2022) additionally make use of molecules’ 3D structures to obtain more informative representations. GraphMVP proposes a generative and a contrastive 3D pre-training task. The generative task can incorporate the information of multiple molecular conformers and adding it to the contrastive pre-training improves downstream performance. Our work differs from GraphMVP in multiple ways. 3D Infomax does not require an additional generative pre-training task to leverage multiple 3D conformers. Instead, we directly include this information in a new contrastive loss formulation. Next, we use multiple 3D datasets for pre-training that belong to different chemical spaces, which allow us to demonstrate the ability to effectively transfer representations between them. Moreover, our evaluation includes quantum mechanical tasks, and we find that the possible improvements in this domain are much larger than for non-quantum properties.

4. 3D Infomax

To achieve our goal of having a 2D GNN that is able to reason about 3D geometry from only 2D inputs, we pre-train our model using contrastive learning. Figure 2 visualizes our

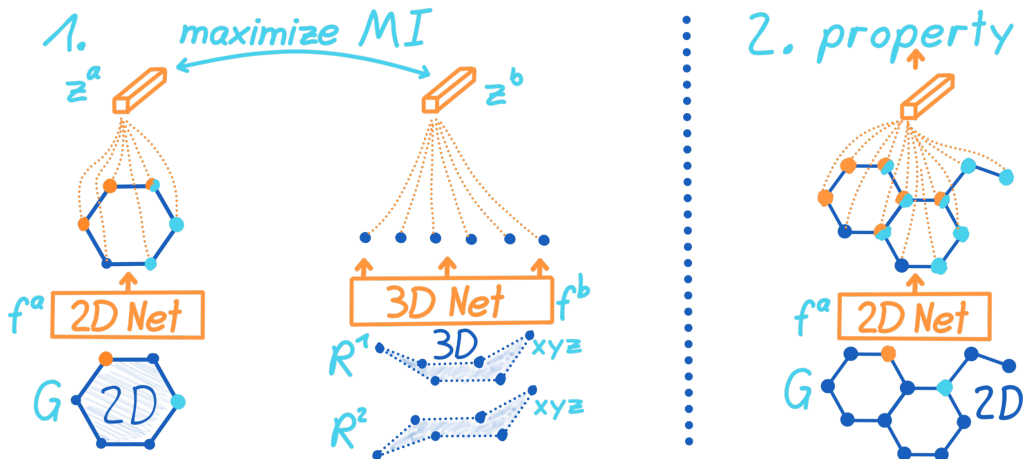


Figure 2. We first pre-train a 2D network f^a by maximizing the mutual information (MI) between its representation z^a of a molecular graph G and a 3D representation z^b produced from the molecules’ conformers R^j . In step 2, the weights of f^a are transferred and fine-tuned to predict properties.

method which maximizes the mutual information between a 2D GNN using the 2D molecular graph and a 3D GNN using the associated 3D conformers. After pre-training, we transfer the weights and fine-tune them on property prediction tasks. During fine-tuning, the GNN’s produced 3D information can be used to improve predictions.

3D Infomax uses two different models, as visualized in Figure 2. Firstly, the model that should be pre-trained which we call *2D network* f^a since its inputs are 2D molecular graphs $G = (\mathcal{V}, \mathcal{E})$ with atoms \mathcal{V} and bonds \mathcal{E} from which it produces a representation $f^a(G) = z^a \in \mathbb{R}^{d_z}$. This can be any GNN that one chooses for the downstream task.

Secondly, the *3D network* f^b which encodes the atoms’ 3D coordinates $R = \{r_v\}_{v \in \mathcal{V}}$ in a 3D representation $f^b(R) = z^b \in \mathbb{R}^{d_z}$. Our pre-training can also be understood from a contrastive distillation (Tian et al., 2019) perspective where the student 2D network learns from the teacher 3D network to produce 3D information.

4.1. Contrastive Framework

To teach the 2D network f^a to produce 3D information from the 2D graph inputs, we maximize the mutual information between the latent 2D representations z^a and 3D representations z^b . Intuitively, we wish to maximize the agreement between z^a and z^b if they are derived from the same molecule. For this purpose, we use contrastive learning (visualized in Figure 3). We consider a batch of N molecular graphs $\{G_i\}_{i \in \{1 \dots N\}}$ with their atom coordinates $\{R_i\}_{i \in \{1 \dots N\}}$ from which the networks produce multiple representations z_i^a and z_i^b .

The first objective of contrastive learning is to maximize the representations’ similarity if they are a positive pair,

meaning that they come from the same molecule (same index i). The second objective is to enforce dissimilarity between negative pairs z_i^a and z_k^b where $i \neq k$, i.e., the 2D and 3D representations in the batch should be dissimilar if they come from different molecules. These objectives are captured in the popular NTXent loss (Chen et al., 2020b) and we use a similar loss to jointly optimize our models:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \left[\log \frac{e^{\text{sim}(z_i^a, z_i^b)/\tau}}{\sum_{k=1, k \neq i}^N e^{\text{sim}(z_i^a, z_k^b)/\tau}} \right] \quad (1)$$

where $\text{sim}(z^a, z^b) = z^a \cdot z^b / (\|z^a\| \|z^b\|)$ is the cosine similarity and τ is a temperature parameter which can be seen as weight for the most similar negative pair. While different combinations of contrastive losses and SSL are possible to learn a joint embedding space between 2D and 3D representations, we found the above loss to perform best. Other methods (Barlow Twins (Zbontar et al., 2021), BYOL (Grill et al., 2020), VICReg (Bardet et al., 2021)) are explored in Appendix C.5.

4.2. Using Multiple Conformers

For most molecules, there are multiple low-energy stable conformers. Instead of only using the most probable conformer (with the lowest energy), we found that leveraging structural information from multiple conformers provides significant benefits. To achieve this, we now consider the c highest probability conformers $\{R_i^j\}_{j \in \{1 \dots c\}}$ of the i -th molecule. If there are fewer than c conformers for a molecule, the lowest energy conformer is repeated. Our choice for the following approach is justified by its good trade-off between simplicity and performance in the comparisons with other possible methods in Appendix C.3.

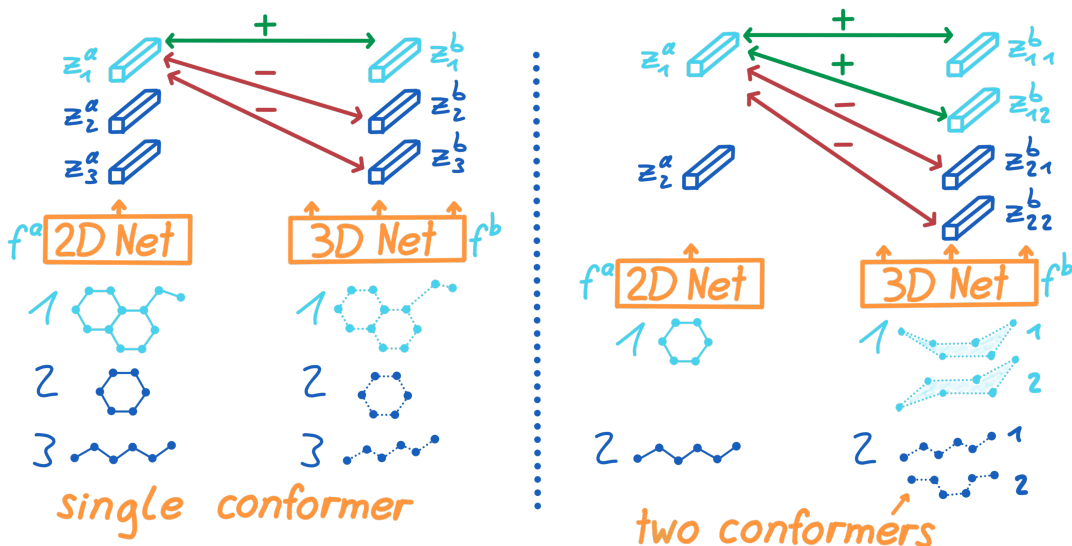


Figure 3. The **single conformer** example shows a batch of three molecular graphs as input to the 2D network with the corresponding three conformers as input to the 3D model. During 3D pre-training, the contrastive loss \mathcal{L} enforces high similarity between latent representations that come from the same molecule (green arrows) while encouraging dissimilarity otherwise (red arrows). This is depicted for the first molecule, but the same is calculated for the second and third. The final loss is the average. The **multiple conformer** example on the right shows two conformers per molecule $c = 2$, and the loss is adjusted to treat all of them as positive pairs if they come from the same molecule and as negative pairs otherwise. Our loss $\mathcal{L}^{multi3D}$ achieves this.

For every molecule the 3D network now takes all conformers as input and produces their latent 3D representations $\{z_{i,j}^b\}_{j \in \{1 \dots c\}}$. The objective is to maximize the similarity between z_i^a and all conformer representations $z_{i,j}^b$ that stem from the same molecule (see Figure 3). As such, we modify our loss to sum over the similarities of all conformers to obtain the final loss:

$$\mathcal{L}^{multi3D} = -\frac{1}{N} \sum_{i=1}^N \left[\log \frac{\sum_{j=1}^c e^{\text{sim}(z_i^a, z_{i,j}^b)/\tau}}{\sum_{k \neq i}^N \sum_{j=1}^c e^{\text{sim}(z_i^a, z_{k,j}^b)/\tau}} \right]. \quad (2)$$

4.3. 3D Network

The 3D network takes as input the coordinates of the atoms as a 3D point cloud and has to produce an SE(3) invariant representation vector z^b that encodes as much information as possible about the 3D structure. Note that it does not have access to the 2D information such as atom or bond features; otherwise, the empirical estimate of the mutual information could be increased by both networks encoding this information instead of the desired 3D structure.

Our concrete architecture encodes the 3D information given by the pairwise Euclidean distances of all atoms. This representation uniquely defines all relative atom positions and is invariant to translation and rotation, as desired. However, like the representations used by many other 3D graph

processing architectures (Klicpera et al., 2020b; Satorras et al., 2021), it is also invariant to reflection, which prevents it from distinguishing chiral molecules. Using all pairwise distances also means that the model’s complexity is quadratic in the number of atoms, which is feasible for drug-like molecules.

The pairwise distances d_{uv} between atoms u and v are first mapped to a higher dimensional space using sine and cosine functions with high frequencies. Recent work provides theoretical (Rahaman et al., 2019) and empirical (Tancik et al., 2020) evidence that this enables deep networks to better fit data with high-frequency variations. Such high-frequency variation is present in the 3D data we wish to encode since differences in bond lengths are often small. Further, the bond distances are likely the most important since close-by atoms typically have the most relevant interactions. As such, we use the following mapping $\gamma : \mathbb{R} \mapsto \mathbb{R}^{2F+1}$ with the number of frequencies F set to 4:

$$\gamma(d_{uv}) = \left(d_{uv}, \sin\left(\frac{d_{uv}}{2^0}\right), \cos\left(d_{uv}/2^0\right), \dots, \sin\left(\frac{d_{uv}}{2^{F-1}}\right), \cos\left(\frac{d_{uv}}{2^{F-1}}\right) \right). \quad (3)$$

The further components can be seen as an MPNN (Gilmer et al., 2017) operating on the fully connected graph of a molecule with the encoded distances as edge features and a constant learned vector as node features. The message

passing layers iteratively encode the 3D information into the node features, which are pooled to produce the 3D representation z^b . The differences to standard MPNNs are detailed in Appendix A. Instead of the presented architecture, a 3D GNN such as SMP (Liu et al., 2021) operating on learned node embeddings could also be used. We validate the improvements provided by our architecture and γ encoding with a set of ablation experiments reported in Appendix C.2.

5. Experiments

5.1. Pre-training Baselines

Distance Predictor. A simpler method to use available 3D structures to pre-train a GNN instead of 3D Infomax is to learn to directly predict all atom distances in the lowest energy conformer. To predict the distance between node v and u , we concatenate their representations $h_u, h_v \in \mathbb{R}^{d_h}$ that were produced by the GNN and feed them to a multi-layer perceptron (MLP) that produces a single scalar $U : \mathbb{R}^{2d_h} \mapsto \mathbb{R}$. The distance prediction $dist_{uv}$ is then given by

$$dist_{uv} = \text{softplus}(U(h_v \parallel h_u) + U(h_u \parallel h_v)) \quad (4)$$

where \parallel denotes concatenation and $\text{softplus}(x) = \log(1 + e^x)$. The node representations are concatenated in both orders and fed to the MLP to ensure that the function is symmetric. The pre-training loss f is the mean squared error between the predicted and true distances.

Conformer Generation. GeoMol (Ganea et al., 2021) is the state-of-the-art deep learning method for generating molecular conformations. It is a generative model that produces a distribution of likely 3D structures for a single molecule, thus capturing the information of multiple conformers. Their architecture employs a GNN whose node representations are used to obtain the final distribution. To use their model as a baseline, we use their training process as a pre-training task and then extract the GNN and fine-tune it on the different downstream tasks.

GraphCL. We compare against the conventional augmentation-based pre-training method GraphCL (You et al., 2020) with the settings of JOAO (You et al., 2021) since it outperformed other SSL approaches for multiple molecular tasks. It uses a common self-supervised objective in which the model has to learn to produce representations that are invariant to augmentations. We use randomly dropping nodes with a ratio of 0.2 on both branches of the SSL setup since JOAO found this combination of augmentations to work particularly well for molecules.

5.2. Data and Evaluation Protocol

To evaluate our method, we pre-train using subsets of 3D molecular structure datasets. We then fine-tune the models

to predict properties of 2D datasets or of 3D datasets while ignoring their 3D information. It is clear that using a 3D GNN and the explicit 3D information of these datasets that was calculated with expensive quantum simulations would yield the best performance. However, we only use the 2D molecular graphs of the fine-tuning data to simulate the scenario when 3D structures are not available. This allows us to compare the improvements through the implicit 3D information of 3D Infomax pre-trained GNN with the maximum possible improvements represented by using the explicit ground truth 3D conformers with a 3D GNN.

The concrete 3D datasets of which we use subsets for pre-training are: **QM9** (Ramakrishnan et al., 2014) which contains 134k small molecules (18 atoms on average) with a single conformer, **GEOM-Drugs** (Axelrod & Gomez-Bombarelli, 2020) with 304k molecules and **QMugs** (Isert et al., 2021) with 665k. GEOM-Drugs and QMugs, both consist of larger drug-like molecules (44.4 and 30.6 atoms on average) with multiple conformers.

For fine-tuning, we predict ten quantum properties of subsets of QM9 and GEOM-Drugs, for which we remove the 3D information. These subsets never have molecules that are contained in the pre-training data, even if the dataset is used to draw pre-training and fine-tuning data from. Additionally, we report results on non-quantum properties of ten OGB (Hu et al., 2020a) datasets in Appendix C.1. The atom and bond features we use for each 2D molecular graph are the same as used in OGB, and further details about all used data are in Appendix B.2.

We choose PNA (Corso et al., 2020) as the GNN to pre-train due to its simplicity and state-of-the-art performance for molecular tasks. The reported confidence intervals are one standard deviation calculated from six random weight initializations, unless stated otherwise. All baselines we compare with use the same GNN as our 3D Infomax method and all experimental settings are detailed in Appendix B. Code to 3D pre-train a GNN, to generate molecular fingerprint embeddings, or to reproduce results is available at <https://github.com/HannesStark/3DInfomax>.

5.3. Quantum Mechanical Properties

Pre-training setup. We use 3D Infomax to pre-train three different instances of PNA (1) on 50k molecules from QM9 using a single conformer, (2) on 140k of GEOM-Drugs with 5 conformers and, (3) on 620k of QMugs using 3 conformers. For comparison, we use two different conventional 2D pre-training methods. These are GraphCL (You et al., 2020) as described in Section 5.1 and pre-training by predicting the Gibbs free energy of GEOM-Drugs’ pre-training subset (labeled *ProPred*). All pre-training methods use a batch size of 500.

Table 1. MAE for QM9’s properties. **3D Infomax** is tested with three different pre-training datasets and **GraphCL** uses a two times larger subset of GEOM-Drugs. **True 3D SMP** is a 3D GNN using ground truth 3D coordinates (hidden from other methods). Details on confidence intervals are in Appendix B. Colors indicate **improvement** (lower MAE) or **worse** performance compared to the randomly initialized (**Rand Init**) model.

TARGET	RAND INIT	PRE-TRAINING BASELINES				OUR 3D INFOMAX			RDKit SMP	True 3D SMP
		GRAPHCL	PROPRED	DISPRED	CONFGEN	QM9	DRUGS	QMUGS		
μ	0.4133 \pm 0.003	0.3937	0.3975	0.4626	0.3940	0.3507	0.3512	0.3668	0.4344	0.0726
α	0.3972 \pm 0.014	0.3295	0.3732	0.3570	0.4219	0.3268	0.2959	0.2807	0.3020	0.1542
HOMO	82.10 \pm 0.33	79.57	93.11	80.58	79.75	68.96	70.78	70.77	82.51	56.19
LUMO	85.72 \pm 1.62	80.81	99.84	84.93	79.16	69.51	71.38	78.10	80.36	43.58
GAP	123.08 \pm 3.98	120.08	131.99	116.21	110.72	101.71	102.59	103.85	114.24	85.10
R2	22.14 \pm 0.21	21.84	29.21	29.23	20.86	17.39	18.96	18.00	22.63	1.51
ZPVE	15.08 \pm 2.83	12.39	11.17	25.91	21.10	7.966	9.677	12.06	5.18	2.69
c_v	0.1670 \pm 0.004	0.1422	0.1795	0.1587	0.1555	0.1306	0.1409	0.1208	0.1419	0.0498

Fine-tuning setup. After pre-training, the models are fine-tuned on 50k molecules from QM9 (in Table 1) or 140k from GEOM-Drugs (in Table 2) that have no overlap with the molecules from the pre-training data. On the same molecules, we also train PNA with random weight initialization (labeled *Rand Init*) to compare how much the downstream performance is improved by the different pre-training methods.

Explicit 3D baseline and ground truth comparison setup.

We additionally train and test the 3D GNN SMP (Liu et al., 2021) with 3D coordinates generated by RDKit’s ETKDG algorithm, (Landrum, 2016) which can be done in a fast manner (labeled *RDKit SMP*). Using conformers generated by the state-of-the-art learned method, GeoMol (Ganea et al., 2021) always performed worse (Appendix C.7). Lastly, we evaluate SMP using the accurate ground truth 3D conformers of QM9 which were computed with time-consuming simulations that would be infeasible for many real-world applications. These structures are not available to the other methods.

3D Infomax for QM9 vs. conventional pre-training. Table 1 shows that 3D Infomax pre-training leads to large improvements over the randomly initialized baseline and over GraphCL with all three pre-training datasets. After 3D pre-training on one half of QM9, the average decrease in MAE is 22%. Comparing 3D Infomax on GEOM-Drugs with GraphCL shows that even though the latter is pre-trained on two times as many molecules from the same dataset, 3D pre-training is always better by a large margin.

Generalization from pre-training to fine-tuning. Pre-training with the disjoint half of QM9 performs best since it shares the molecular space of the test set. Nevertheless, the learned representations also generalize well: pre-training on GEOM-Drugs and QMugs leads to improvements of 19% and 18% respectively, even though QM9 contains much smaller molecules with on average 18 atoms compared to the 44.4 atoms for the drug-like molecules of GEOM-Drugs.

Comparison with Ground Truth Conformers. 3D Infomax’ large improvements have to be compared to methods that also do not use explicit ground truth 3D information and only operate on the 2D molecular graph for the molecules of which properties should be predicted. However, for the quantum property experiment datasets we employ, high-quality ground truth 3D information was calculated with expensive quantum simulations and we are thus able to compare with 3D GNNs that use this additional input for property prediction.

The results of SMP in Table 5.3 show that for some properties, the MAE of 3D Infomax with its implicit 3D information is still higher than what is possible when using explicit ground truth conformers (which are time-consuming and expensive to obtain). One likely contributing factor for this is that QM9’s properties are conformer-specific. There might be a maximum accuracy that can be achieved if only the molecule is known and not for which conformer the property should be predicted. Nevertheless, this performance gap suggests that there is still room for improvement.

3D pre-training Baselines. 3D pre-training by directly predicting 3D quantities is simpler than 3D Infomax and would be preferable in case of similar gains. Therefore, we compare with the baselines in Section 5.1 using the same 140k molecules of GEOM-Drugs for all 3D pre-training methods. *DisPred* refers to predicting all atom distances of the highest probability conformer, and *ConfGen* means pre-training by predicting up to 10 conformers. Table 5.3 shows that 3D Infomax pre-training is always superior to the 3D pre-training baselines and is the only method to not suffer from negative transfer (Pan & Yang, 2010).

3D Infomax for GEOM-Drugs. Table 2 further confirms that 3D Infomax substantially improves quantum property predictions and generalizes out-of-distribution. Our method outperforms GraphCL, even though GraphCL also sees the fine-tuning molecules during pre-training. Moreover, we observe strong generalization when pre-training with QM9

Table 2. The MAE for predicting GEOM-Drugs’ properties. **3D Infomax** compared with **GraphCL** and no pre-training.

METHOD	GIBBS	$\langle E \rangle$
RAND INIT	.2035	.1026
GRAPHCL	.1941	.0995
3D INFOMAX QM9	.1852	.0968
3D INFOMAX DRUGS	.1811	.0952
3D INFOMAX QMUGS	.1835	.0965

and fine-tuning on GEOM-Drugs. In this case, the pre-training data only contains the elements C, H, N, O, and F while the target data contains eleven additional elements that are unseen during pre-training.

Results interpretation. Such consistent and out-of-distribution improvements can be explained by the type of information captured with 3D Infomax. Learning to reason about molecular geometry and its impact does not depend on the data’s molecular space. Therefore it is not necessary to have a high similarity between the molecules during pre-training and fine-tuning.

Another advantage of 3D Infomax is its comparably fast convergence. Pre-training on 620k molecules of QMugs with 3 conformers takes 12 hours, compared to 71 hours for GraphCL on 280k molecules of GEOM-Drugs.

5.4. Number of Conformers and pre-training Molecules

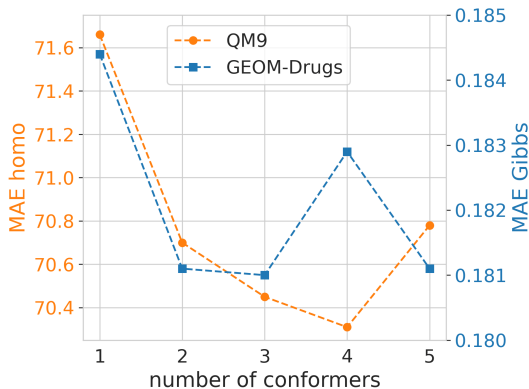


Figure 4. The MAE for the **QM9’s homo** and **GEOM-Drugs’ Gibbs** property when varying the number of GEOM-Drugs’ conformers used during pre-training.

Figure 4 highlights the benefit of using more than a single conformer. However, the marginal gain reduces as higher energy conformers are added and beyond a certain point (around three conformers), the reduced focus on the most likely conformers worsens the downstream performance. This is in line with the observation that, on average, three conformers are enough to cover 70% of the cumulative Boltzmann weight for GEOM-Drugs. Additionally, experi-

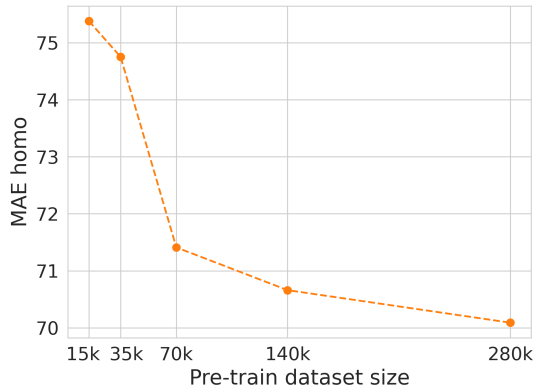


Figure 5. The MAE when using different numbers of molecules of GEOM-Drugs during pre-training.

ments in Appendix C.3 show that using multiple conformers is essential when pre-training with QMugs: the MAE for QM9’s homo property is 82.57 with a single conformer while it improves to 70.77 when using three.

In Figure 5 we can observe the performance improving as the size of the pre-training dataset increases. However, the returns are diminishing, and we cannot claim that even larger pre-training datasets are likely to drastically improve performance.

6. Conclusion

We presented a pre-training strategy, 3D Infomax, that teaches a GNN to produce latent 3D and quantum information from 2D molecular graphs. This can later be used during fine-tuning to improve molecular property predictions while retaining the inference speed of a standard GNN operating on 2D molecular graphs. We found consistently large improvements (~22%) for quantum properties, overshadowing the gains possible with conventional SSL methods. The embedded 3D knowledge can be transferred across highly different types of molecules (e.g., from molecules with an average of 18 atoms to drug-like molecules with 44.4 atoms) since the representations capture a principled form of information that is known to be useful for several molecular tasks. Lastly, we observed that using multiple molecular conformers during pre-training provides valuable additional information to further improve downstream property predictions.

We provide open-source access to our method at <https://github.com/HannesStark/3DInfomax> with a simple setup to 3D pre-train a GNN and the additional possibility to generate molecular fingerprint embeddings that carry latent 3D information from a list of SMILES.

7. Acknowledgments

The authors express their gratitude to Simon Axelrod (GEOM Dataset and molecular physics advice), Limei Wang and Meng Liu (Spherical Message Passing), Adrien Bardes and Samuel Lavoie (SSL, VICReg, and BYOL), Octavian Ganea and Lagnajit Pattanaik (molecular geometry, GeoMol), Christopher Scarveils (part of the project in its beginnings), Bertrand Charpentier and Aleksandar Bojchevski (general help) as well as Zekarias Kefato, Grégoire Mialon, Kenneth Atz, Johannes Klicpera, Minghao Xu, Mozhi Zhang, Shantanu Thakoor, Minkai Xu, Shitong Luo, Yuning You, and Prannay Khosla for insightful discussions.

CD acknowledges support by the Bundesministerium für Bildung und Forschung (BMBF), project number 01IS17049.

References

- Axelrod, S. and Gomez-Bombarelli, R. Geom: Energy-annotated molecular conformations for property prediction and molecular generation. *arXiv preprint arXiv:2006.05531*, 2020.
- Axelrod, S. and Gómez-Bombarelli, R. Molecular machine learning with conformer ensembles. *CoRR*, abs/2012.08452, 2020.
- Bardes, A., Ponce, J., and LeCun, Y. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *CoRR*, abs/2105.04906, 2021.
- Belghazi, M. I., Baratin, A., Rajeswar, S., Ozair, S., Bengio, Y., Hjelm, R. D., and Courville, A. C. Mutual information neural estimation. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 530–539, 2018.
- Bemis, G. W. and Murcko, M. A. The properties of known drugs. 1. molecular frameworks. *Journal of Medicinal Chemistry*, 39(15):2887–2893, 1996.
- Brockschmidt, M. Gnn-film: Graph neural networks with feature-wise linear modulation. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1144–1152. PMLR, 2020.
- Bronstein, M. M., Bruna, J., Cohen, T., and Velickovic, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *CoRR*, abs/2104.13478, 2021.
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., and Joulin, A. Unsupervised learning of visual features by contrasting cluster assignments. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Chen, P., Liu, W., Hsieh, C.-Y., Chen, G., and Heng, P. A. Utilizing edge features in graph neural networks via variational information maximization, 2020a.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. E. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1597–1607. PMLR, 2020b.
- Chen, X. and He, K. Exploring simple siamese representation learning. *CoRR*, abs/2011.10566, 2020.
- Coley, C., Jin, W., Rogers, L., Jamison, T. F., Jaakkola, T. S., Green, W. H., Barzilay, R., and Jensen, K. F. A graph-convolutional neural network model for the prediction of chemical reactivity. *Chem. Sci.*, 10:370–377, 2019.
- Corso, G., Cavalleri, L., Beaini, D., Liò, P., and Veličković, P. Principal neighbourhood aggregation for graph nets. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 13260–13271. Curran Associates, Inc., 2020.
- Dral, P. O. Quantum chemistry in the age of machine learning. *The Journal of Physical Chemistry Letters*, 11(6): 2336–2347, 2020.
- Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- Ganea, O.-E., Pattanaik, L., Coley, C. W., Barzilay, R., Jensen, K. F., Green, W. H., and Jaakkola, T. S. Geomol: Torsional geometric generation of molecular 3d conformer ensembles, 2021.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In Teh, Y. W. (ed.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1263–1272. PMLR, 2017.
- Gorgulla, C., Boeszoermyeni, A., Wang, Z.-F., Fischer, P. D., Coote, P. W., Padmanabha Das, K. M., Malets, Y. S., Radchenko, D. S., Moroz, Y. S., Scott, D. A., Fackeldey,

- K., Hoffmann, M., Iavniuk, I., Wagner, G., and Arthanari, H. An open-source drug discovery platform enables ultralarge virtual screens. *Nature*, 580(7805):663–668, 2020.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.
- Grill, J., Strub, F., Alché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. Á., Guo, Z., Azar, M. G., Piot, B., Kavukcuoglu, K., Munos, R., and Valko, M. Bootstrap your own latent - A new approach to self-supervised learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Grimme, S. Exploration of chemical compound, conformer, and reaction space with meta-dynamics simulations based on tight-binding quantum chemical calculations. *Journal of Chemical Theory and Computation*, 15(5):2847–2862, 2019.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In Teh, Y. W. and Titterton, D. M. (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, volume 9 of *JMLR Proceedings*, pp. 297–304. JMLR.org, 2010.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020a.
- Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V. S., and Leskovec, J. Strategies for pre-training graph neural networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020b.
- Hy, T. S., Trivedi, S., Pan, H., Anderson, B. M., and Kondor, R. Predicting molecular properties with covariant compositional networks. *The Journal of Chemical Physics*, 148(24):241745, 2018.
- Isert, C., Atz, K., Jiménez-Luna, J., and Schneider, G. Qmugs: Quantum mechanical properties of drug-like molecules, 2021.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P., and Hassabis, D. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 8 2021.
- Klicpera, J., Giri, S., Margraf, J. T., and Günnemann, S. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. In *NeurIPS-W*, 2020a.
- Klicpera, J., Groß, J., and Günnemann, S. Directional message passing for molecular graphs. In *International Conference on Learning Representations (ICLR)*, 2020b.
- Klicpera, J., Becker, F., and Günnemann, S. Gemnet: Universal directional graph neural networks for molecules. *NeurIPS*, 2021.
- Landrum, G. Rdkit: Open-source cheminformatics software, 2016.
- Li, J., Cai, D., and He, X. Learning graph-level representation for drug discovery. *CoRR*, abs/1709.03741, 2017.
- Liu, S., Wang, H., Liu, W., Lasenby, J., Guo, H., and Tang, J. Pre-training molecular graph representation with 3d geometry. *10th International Conference on Learning Representations, ICLR 2022*, 2022.
- Liu, Y., Wang, L., Liu, M., Zhang, X., Oztekin, B., and Ji, S. Spherical message passing for 3d graph networks. *CoRR*, abs/2102.05013, 2021.
- Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.*, 22(10):1345–1359, 2010.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch, 2017.
- Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., and Courville, A. On the spectral bias of neural networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of

- Proceedings of Machine Learning Research*, pp. 5301–5310. PMLR, 6 2019.
- Ramakrishnan, R., Dral, P. O., Rupp, M., and von Lilienfeld, O. A. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1(1):140022, 8 2014.
- Satorras, V. G., Hoogeboom, E., and Welling, M. E(n) equivariant graph neural networks. *CoRR*, abs/2102.09844, 2021.
- Schmidt, J., Marques, M. R. G., Botti, S., and Marques, M. A. L. Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials*, 5(1):83, 8 2019.
- Schütt, K., Kindermans, P., Felix, H. E. S., Chmiela, S., Tkatchenko, A., and Müller, K. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 991–1001, 2017.
- Shi, C., Luo, S., Xu, M., and Tang, J. Learning gradient fields for molecular conformation generation. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9558–9568. PMLR, 2021.
- Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C. R., French, S., Carfrae, L. A., Bloom-Ackermann, Z., Tran, V. M., Chiappino-Pepe, A., Badran, A. H., Andrews, I. W., Chory, E. J., Church, G. M., Brown, E. D., Jaakkola, T. S., Barzilay, R., and Collins, J. J. A deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702.e13, 2020.
- Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.
- Tang, B., Kramer, S. T., Fang, M., Qiu, Y., Wu, Z., and Xu, D. A self-attention based message passing neural network for predicting molecular lipophilicity and aqueous solubility. *J. Cheminformatics*, 12(1):15, 2020.
- Tian, Y., Krishnan, D., and Isola, P. Contrastive representation distillation. *CoRR*, abs/1910.10699, 2019. URL <http://arxiv.org/abs/1910.10699>.
- Torng, W. and Altman, R. B. Graph convolutional neural networks for predicting drug-target interactions. *J. Chem. Inf. Model.*, 59(10):4131–4149, 2019.
- Unke, O. T. and Meuwly, M. Physnet: A neural network for predicting energies, forces, dipole moments, and partial charges. *Journal of Chemical Theory and Computation*, 15(6):3678–3693, 2019.
- van den Oord, A., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748, 2018.
- Wang, M., Zheng, D., Ye, Z., Gan, Q., Li, M., Song, X., Zhou, J., Ma, C., Yu, L., Gai, Y., Xiao, T., He, T., Karypis, G., Li, J., and Zhang, Z. Deep graph library: A graph-centric, highly-performant package for graph neural networks. *arXiv preprint arXiv:1909.01315*, 2019.
- Wang, Y., Wang, J., Cao, Z., and Farimani, A. B. Molclr: Molecular contrastive learning of representations via graph neural networks. *CoRR*, abs/2102.10056, 2021.
- Withnall, M., Lindelöf, E., Engkvist, O., and Chen, H. Building attention and edge message passing neural networks for bioactivity and physical-chemical property prediction. *J. Cheminformatics*, 12(1):1, 2020.
- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. S. Moleculenet: A benchmark for molecular machine learning. *CoRR*, abs/1703.00564, 2017.
- Xu, M., Luo, S., Bengio, Y., Peng, J., and Tang, J. Learning neural generative dynamics for molecular conformation generation. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021a.
- Xu, M., Wang, H., Ni, B., Guo, H., and Tang, J. Self-supervised graph-level representation learning with local and global structure. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11548–11558. PMLR, 2021b.
- Yang, K., Swanson, K., Jin, W., Coley, C., Eiden, P., Gao, H., Guzman-Perez, A., Hopper, T., Kelley, B., Mathea, M., Palmer, A., Settels, V., Jaakkola, T., Jensen, K., and Barzilay, R. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling*, 59(8):3370–3388, 2019.
- You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., and Shen, Y. Graph contrastive learning with augmentations. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and

Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

You, Y., Chen, T., Shen, Y., and Wang, Z. Graph contrastive learning automated. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12121–12132. PMLR, 2021.

Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. Barlow twins: Self-supervised learning via redundancy reduction. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12310–12320. PMLR, 2021.

Zhu, J., Xia, Y., Qin, T., Zhou, W., Li, H., and Liu, T. Dual-view molecule pre-training. *CoRR*, abs/2106.10234, 2021.

A. Further Explanations

3D Network Details The l -th layer of the 3D network takes two sets as input. First, $n^2 - n$ edge representations $\{d_{uv}^l \in \mathbb{R}^{d_d} \mid u, v \in \mathcal{V} \wedge u \neq v\}$ (the edges of a complete graph without self-loops). In the first layer they are given by the encoded distances fed through an initial feed-forward network $U_{init} : \mathbb{R}^{2F+1} \mapsto \mathbb{R}^{d_d}$ which projects them to the hidden dimension of the edges $d_{uv}^0 = U_{init}(\gamma(d_{uv}))$. The second input is a set of n atom representations $\{h_1^l, \dots, h_n^l\}$ with dimensionality \mathbb{R}^{d_h} . In the first layer, the atom representations are all set to the same learned vector that is initialized with a standard normal. With \parallel meaning concatenation, every layer updates the edge and atom representations and iteratively encodes 3D information into them as follows:

$$m_{uv} = U_{edge}([h_u^l \parallel h_v^l \parallel d_{uv}^l]) \quad (5)$$

$$d_{uv}^{l+1} = d_{uv}^l + m_{uv} \quad (6)$$

$$h_u^{l+1} = U_h([h_u \parallel \sum_{\substack{v=1 \\ v \neq u}}^n m_{uv} * \sigma(U_{softedge}(m_{uv}))]). \quad (7)$$

The layer is parameterized by three MLPs where the first one updates the edges $U_{edge} : \mathbb{R}^{2d_h+d_d} \mapsto \mathbb{R}^{d_d}$. The second one updates the atom representations $U_h : \mathbb{R}^{d_h+d_d} \mapsto \mathbb{R}^{d_h}$. The third one $U_{softedge} : \mathbb{R}^{d_d} \mapsto \mathbb{R}$ is followed by the logistic sigmoid function to create a value between 0 and 1 that can be seen as a soft edge weight telling us how probable an edge is for each message m_{uv} as it is done by Satorras et al. (2021).

To produce the final 3D representation z^b , all atom representations are aggregated by concatenating their mean, their maximum, and their standard deviation and feeding this through a final feed-forward network $U : \mathbb{R}^{3d_h} \mapsto \mathbb{R}^{d_z}$.




Figure 6. Depiction of two groups of molecules where all the molecules in the top row have the same Bemis-Murcko scaffold (Bemis & Murcko, 1996) which is different from the scaffold to which the two molecules in the bottom row belong. We can easily obtain the scaffold of a molecule using RDKit (Landrum, 2016).

B. Experimental Details

B.1. Parameter Details

The hyperparameters for SMP are taken from the official repository¹ where Liu et al. (2021) provide their code, and we predict *gap* even though it could be calculated as $|homo - lumo|$. The parameter search space and final parameters for the PNA architecture are specified in Table 3 and those of the 3D network in Table 4.

¹<https://github.com/divelab/DIG>

Pre-training: We use Adam with a start learning rate of 8×10^{-5} and a batch size of 500. The learning rate schedule during pre-training starts with 700 optimization steps of linear warmup followed by the schedule given by the *ReduceLRonPlateau* scheduler by PyTorch² with reduction parameter 0.6, patience 25, and a cooldown of 20.

Fine-tuning quantum mechanical properties: We use Adam with a start learning rate of 7×10^{-5} , weight decay 1×10^{-11} and a batch size of 128. For the learning rate schedule, we first perform warmup as follows. We consider three different sets of learnable parameters: (1) the batch norm parameters, (2) all newly initialized parameters that were not transferred, and (3) all parameters. For these sets, we increase the learning rate in this order from 0 to the start learning rate with linear interpolation. For parameter group one, we warm up for 700 steps, 700 steps for group 2, and 350 steps for group 3. After that we use the schedule given by the *ReduceLRonPlateau* with reduction parameter 0.5, patience 25, and a cooldown of 20.

Fine-tuning non-quantum properties: We use Adam with a start learning rate of 1×10^{-3} and a batch size of 32. The learning rate schedule is the same as for the quantum mechanical properties.

The experiment on the non-quantum properties has different hyperparameters for PNA since the smaller datasets are easily overfitted on with the large architecture we use for the quantum mechanical properties. A smaller PNA yields better performance with the random initialization baseline. Therefore the PNA in these experiments has a hidden dimension of 50 and 3 message passing layers as propagation depth. Apart from that, it is the same as the PNA described in Table 3.

Table 3. Search space for the 2D network PNA through which we searched to obtain a strong baseline performance on the homo property of the QM9 dataset. The parameters were tuned in the order in which they are listed in this table from top to bottom. After this was completed for all parameters, we performed a second round of tuning for a subset of them. The final parameters are marked in **bold**.

PARAMETER	SEARCH SPACE
PROPAGATION DEPTH	[4, 5, 6, 7]
HIDDEN DIMENSION	[40, 50, 75, 90, 100, 150, 200 , 300]
MESSAGE MLP LAYERS	[1, 2 , 3]
UPDATE MLP LAYERS	[1 , 2, 3]
AGGREGATORS	[MEAN, MAX, MIN, STD, SUM], [MEAN, MAX, MIN], [MEAN, MAX, SUM], [MEAN, MAX, MIN, STD], [MAX, SUM], [SUM]
SCALERS	[IDENTITY], [IDENTITY , AMPLIFICATION, ATTENUATION]
READOUT AGGREGATORS	[MEAN], [SUM], [MEAN, MAX, SUM], [MEAN, MAX, MIN, SUM]
DROPOUT	[0 , 0.05, 0.1, 0.2]
BATCHNORM AFTER MLPs	TRUE /FALSE
BATCHNORM IN MLPs	TRUE /FALSE
READOUT MLP LAYERS	[1, 2 , 3]
BATCHNORM MOMENTUM	[0.1 , 0.9, 0.93]

Table 4. Search space for the 3D network *Net3D* through which we searched to obtain a strong baseline performance on the homo property of the QM9 dataset and we considered the size of the network where parameters leading to less memory use are preferred. The parameters were tuned in the order in which they are listed in this table from top to bottom. After this was completed for all parameters, we performed a second round of tuning for a subset of them. The final parameters are marked in **bold**.

PARAMETER	SEARCH SPACE
PROPAGATION DEPTH	[1 , 3, 4, 5]
HIDDEN DIMENSION	[10, 20 , 40, 60, 80, 100]
F USED IN $\gamma : \mathbb{R} \mapsto \mathbb{R}^{2F+1}$	[0, 3, 4 , 8, 10, 50]
MESSAGE MLP LAYERS	[1, 2 , 3]
UPDATE MLP LAYERS	[1 , 2, 3]
READOUT AGGREGATORS	[MEAN], [SUM], [MEAN, MAX, MIN], [MEAN, MAX, MIN, SUM]
DROPOUT	[0 , 0.05, 0.1, 0.2, 0.5]
BATCHNORM AFTER MLPs	TRUE /FALSE
READOUT MLP LAYERS	[1, 2 , 3]
BATCHNORM MOMENTUM	[0.1, 0.9, 0.93]

²https://pytorch.org/docs/stable/generated/torch.optim.lr_scheduler.ReduceLRonPlateau.html

B.2. Data Details

We use three datasets containing 3D information for pre-training with diversity in molecule size and the number of molecules, as can be seen in Table 5. The pre-training datasets are:

1. **QM9**³ (Ramakrishnan et al., 2014) contains 134k stable small organic molecules of 5 elements (CHONF). Every molecule has the 3D coordinates of one low-energy conformer and is annotated with 12 quantum mechanical properties as regression targets. The molecules are considered very small, with at most 9 heavy atoms.
2. **GEOM-Drugs**⁴ (Axelrod & Gomez-Bombarelli, 2020) consists of 304k realistically-sized biologically and pharmacologically relevant molecules of 16 elements, annotated with multiple 3D conformers, the ensemble Gibbs free energy, and the ensemble energy as regression targets. For the average molecule, 70% of the Boltzmann weight is captured by just three conformers as can be seen in Figure B.2 where we also provide a histogram for the number of molecules that have a certain amount of conformers in Figure B.2. The conformers are generated using CREST (Grimme, 2019).
3. **QMugs**⁵ (Isert et al., 2021) has 665k drug-like molecules with three diverse conformers each and multiple conformer specific quantum mechanical properties as regression tasks. The conformers are generated using CREST (Grimme, 2019).

For fine-tuning, we use a variety of datasets that cover a wide range of domains and applications. The molecular properties are relevant for quantum mechanics, physical chemistry, biophysics, and physiology such that we can obtain a good estimate of how valuable our 3D pre-training is for each domain. For quantum mechanical properties, which are often specific to a conformer, it is clear that 3D information is important and there has been a lot of evidence that learned methods highly benefit from its use (Klicpera et al., 2020b;a; Liu et al., 2021; Schütt et al., 2017). For these properties, the interest is in how much our method can leverage this information and transfer it to molecules where no 3D geometry is available.

Meanwhile, for biological or physiological properties such as blood-brain barrier penetration, it is not as clear if improvements from 3D information are to be expected. As such, this question needs to be answered next to how much of the benefits 3D pre-training recovers. For this purpose, we use the following molecular graph datasets, which are mainly taken from MoleculeNet (Wu et al., 2017) and we use the scaffold splits⁶ with an 80/10/10 split ratio provided by OGB Hu et al. (2020a). The fine-tuning datasets are:

1. **QM9 and GEOM-Drugs:** On these 3D datasets we also fine-tune and evaluate the quantum mechanical properties of one half of the datasets with a random split. This is done after either pre-training on another 3D dataset (generalization), or after pre-training on the other half of the same dataset (in distribution).
2. **ESOL:** 1128 common organic small molecules with water solubility data (log solubility in mols per liter).
3. **Lipo:** Experimental data for the octanol/water distribution coefficient of 4200 molecules.
4. **FreeSolv:** The hydration free energy of 642 molecules in water.
5. **HIV:** 41k molecules with binary labels for HIV virus replication inhibition.
6. **BACE:** Binary labels of binding results for inhibitors of human β -secretase 1 for 1512 molecules.
7. **BBBP:** 2039 molecules with binary labels for blood-brain barrier penetration.
8. **Tox21:** 7831 molecules with binary labels of their toxic for 12 different targets.
9. **ToxCast:** 8576 molecules with binary labels of toxicity experiment outcomes with 617 targets.
10. **SIDER:** 1427 approved drugs with 27 different side effect groups and the task is to predict whether the drug is in the side effect group.
11. **ClinTox:** 1477 drugs with two binary annotations where the first is to predict toxicity in clinical trials and the second is the FDA approval status.

The reason why *muv* and *pcha* are the only datasets from the OGB benchmark suite which we omit is their larger size.

³https://github.com/klicperaajo/dimenet/blob/master/data/qm9_eV.npz

⁴<https://github.com/learningmatter-mit/geom>

⁵<https://www.research-collection.ethz.ch/handle/20.500.11850/482129>

⁶<https://ogb.stanford.edu/docs/graphprop>

Table 5. Statistics of the used datasets. In the upper section are datasets with 3D information, which we use for pre-training, and the datasets in the bottom section do not contain additional 3D annotations.

DATASET	#MOLECULES	AVG. #ATOMS	AVG. #BONDS	SPLIT
QM9	130 831	18.0	18.6	RANDOM
GEOM-DRUGS	304 293	44.4	46.4	RANDOM
QMUGS	665 911	30.6	33.4	RANDOM
ESOL	1128	13.3	13.7	SCAFFOLD
LIPO	4200	27.0	29.5	SCAFFOLD
FREESOLV	642	8.7	8.4	SCAFFOLD
BACE	1512	34.1	36.9	SCAFFOLD
BBBP	2039	24.1	26.0	SCAFFOLD
HIV	41 127	25.5	27.5	SCAFFOLD
TOX21	7831	18.6	19.3	SCAFFOLD
TOXCAST	8576	18.8	19.3	SCAFFOLD
CLINTOX	1477	26.2	27.9	SCAFFOLD
SIDER	1427	33.6	35.4	SCAFFOLD

B.3. Implementation

Code to 3D pre-train a GNN or to reproduce results is available at <https://github.com/HannesStark/3DInfomax>. All experiments were implemented in *PyTorch* (Paszke et al., 2017) using the deep learning libraries for processing graphs *Pytorch Geometric* (Fey & Lenssen, 2019) and *Deep Graph Library* (Wang et al., 2019). The code we use for SMP (Liu et al., 2021) is under the GNU General Public License v3.0 and we use their implementation after discussing it with the first author of the paper and under the consideration that their project welcomed our contributions to their library.

The experiments were conducted on two different machines while the same system was always used in direct comparisons. The first machine has an AMD Ryzen 1700 CPU @ 3.70Ghz, 16GB of RAM, and an Nvidia GTX 1060 GPU with 6GB vRAM. The second system contains two Intel Xeon Gold 6248 CPUs @ 2.50GHz each with 20/40 cores, 400GB of RAM, and four Quadro RTX 8000 GPUs with 46GB vRAM of which only a single one was used for each experiment. All mentions and of training time refer to the second system.

B.4. Units and Meaning of Quantum Properties

For the GEOM-Drugs dataset, all reported numbers have the unit kcal/mol, Gibbs refers to the ensemble Gibbs free energy, and $\langle E \rangle$ to the ensemble energy.

Table 6. Units and description of quantum mechanical properties of the QM9 dataset.

PROPERTY	UNIT	DESCRIPTION
μ	DEBYE	DIPOLE MOMENT
α	$Bohr^3$	ISOTROPIC POLARIZABILITY
<i>homo</i>	MEV	ENERGY OF HIGHEST OCCUPIED MOLECULAR ORBITAL (HOMO)
<i>lumo</i>	MEV	ENERGY OF LOWEST OCCUPIED MOLECULAR ORBITAL (LUMO)
<i>gap</i>	MEV	GAP, DIFFERENCE BETWEEN LUMO AND HOMO
<i>r2</i>	$Bohr^2$	ELECTRONIC SPATIAL EXTENT
<i>ZPVE</i>	MEV	ZERO POINT VIBRATIONAL ENERGY
<i>c_v</i>	$\frac{cal}{molK}$	HEAT CAPACITY AT 298.15 K

B.5. Confidence Interval Details

All the specified confidence intervals in our work are standard deviations calculated from different weight initializations using the seeds [1, 2, 3, 4, 5, 6] or [1, 2, 3, 4]. The following tables provide additional confidence intervals for the results in the main text.

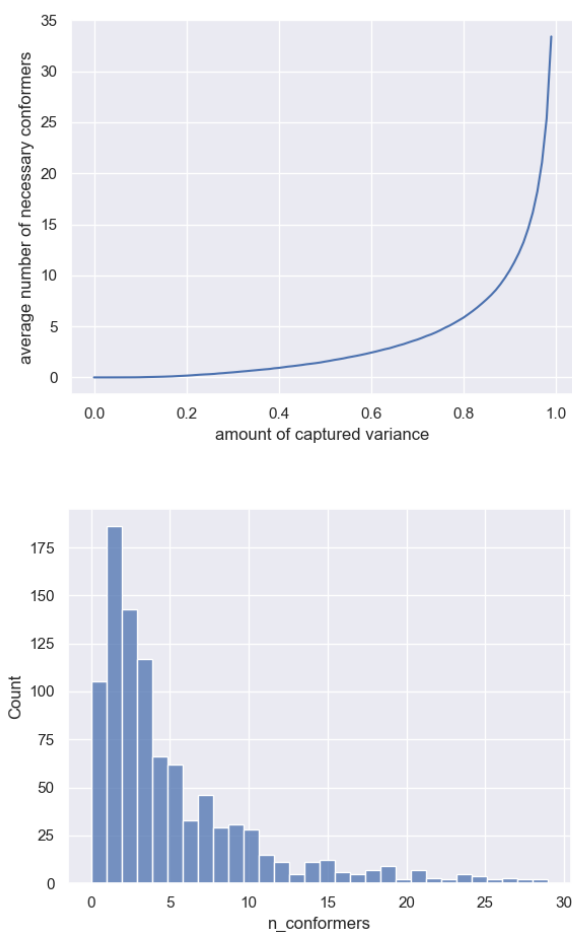


Figure 7. a) The average number of conformers necessary to cover a certain amount of Boltzmann weight in GEOM-Drugs. For a given amount of cumulative Boltzmann weight on the horizontal axis, the vertical axis shows the average number of conformers necessary to pass that threshold. b) Histogram of how many molecules there are in GEOM-Drugs with a certain amount of conformers. The histogram is created for 1000 molecules of GEOM-Drugs.

C. Additional Results

C.1. Non-Quantum Properties

We found that 3D Infomax yields large improvements for predicting quantum properties. For non-quantum properties, there is less empirical evidence that explicit 3D information improves prediction accuracy with first explorations by [Axelrod & Gómez-Bombarelli \(2020\)](#) showing only little improvements for their COVID 19 related predictions. Nevertheless, for tasks such as binding prediction in the *bace* dataset, we would expect it to be helpful, and we compare different methods pre-trained with GEOM-Drugs.

In Table 10, we find that 3D Infomax improved performance for 4 out of 10 OGB datasets. In contrast to the results for quantum mechanical property predictions (Section 5.3), it is not always superior to GraphCL and ConfGen. However, 3D Infomax never decreases performance which can be valuable in practice and make the method worth employing for non-quantum properties as well.

When investigating for which tasks 3D Infomax is useful, we see that abstract tasks such as predicting clinical test outcomes (*clintox*) benefit less. The most significant improvements are rather possible for tasks like predicting solubility and lipophilicity in *esol* and *lipo*. These are more directly related to molecular mechanics and a molecule’s intrinsic properties (e.g., the dipole moment/polarity is important for predicting lipophilicity). They do not depend on how a molecule

Table 7. Additional confidence intervals of our method in Table 1. All standard deviations are calculated from 4 seeds except for the homo property where 6 are used.

TARGET	RAND INIT	OUR 3D INFOMAX		
		QM9	DRUGS	QMUGS
μ	0.4133 \pm 0.003	0.3507 \pm 0.005	0.3512 \pm 0.010	0.3668 \pm 0.004
α	0.3972 \pm 0.014	0.3268 \pm 0.006	0.2959 \pm 0.009	0.2807 \pm 0.012
HOMO	82.10 \pm 0.33	68.96 \pm 0.32	70.78 \pm 0.82	70.77 \pm 0.74
LUMO	85.72 \pm 1.62	69.51 \pm 0.54	71.38 \pm 0.74	78.10 \pm 0.69
GAP	123.08 \pm 3.98	101.71 \pm 2.03	102.59 \pm 3.27	103.85 \pm 1.92
R2	22.14 \pm 0.21	17.39 \pm 0.94	18.96 \pm 0.69	18.00 \pm 0.40
ZPVE	15.08 \pm 2.83	7.966 \pm 1.87	9.677 \pm 1.29	12.06 \pm 2.40
c_v	0.1670 \pm 0.004	0.1306 \pm 0.009	0.1409 \pm 0.016	0.1208 \pm 0.008

Table 8. Additional confidence intervals for Table 2.

METHOD	GIBBS	$\langle E \rangle$
RAND INIT	.2035 \pm 0.0011	.1026 \pm 0.0017
GRAPHCL	.1941	.0995
3D INFOMAX QM9	.1852	.0968
3D INFOMAX DRUGS	.1811	.0952
3D INFOMAX QMUGS	.1835	.0965

Table 9. Additional confidence intervals for 3D pre-training baselines in Table 5.3

METHOD	μ	α	HOMO	LUMO	GAP	R2	ZPVE	c_v
RAND INIT	\pm 0.003	\pm 0.014	\pm 0.33	\pm 1.62	\pm 3.98	\pm 0.21	\pm 2.83	\pm 0.004
3D INFOMAX	\pm 0.010	\pm 0.009	\pm 0.82	\pm 0.74	\pm 3.27	\pm 0.69	\pm 1.29	\pm 0.016

Table 10. Comparison of 3D pre-training baselines and **GraphCL** against **3D Infomax** on various OGB datasets. Shown is either the root mean squared error (RMSE) (lower is better) or the area under the ROC-curve (ROC-AUC) (higher is better). Colors indicate **improvement**, **worse** performance, or **no significant change** compared to the randomly initialized (**Rand Init**) model.

DATASET	METRIC	RAND INIT	DISPRED	CONFGEN	GRAPHCL	3D INFOMAX
ESOL	RMSE \downarrow	0.947 \pm 0.038	0.986 \pm 0.025	0.867 \pm 0.045	0.959 \pm 0.047	0.894 \pm 0.028
LIPO	RMSE \downarrow	0.739 \pm 0.009	0.718 \pm 0.021	0.757 \pm 0.035	0.714 \pm 0.011	0.695 \pm 0.012
FRESOLV	RMSE \downarrow	2.233 \pm 0.261	2.486 \pm 0.222	2.428 \pm 0.155	3.744 \pm 0.292	2.337 \pm 0.227
BACE	ROC-AUC \uparrow	78.13 \pm 1.30	76.51 \pm 1.95	80.02 \pm 1.58	77.18 \pm 4.01	79.42 \pm 1.94
BBBP	ROC-AUC \uparrow	68.27 \pm 1.98	66.06 \pm 1.84	66.16 \pm 2.24	71.06 \pm 2.00	69.10 \pm 1.07
TOX21	ROC-AUC \uparrow	73.88 \pm 0.64	73.87 \pm 0.43	75.24 \pm 1.00	78.92 \pm 0.61	74.46 \pm 0.74
TOXCAST	ROC-AUC \uparrow	63.62 \pm 0.48	61.58 \pm 0.58	64.74 \pm 1.20	64.95 \pm 0.31	64.41 \pm 0.88
CLINTOX	ROC-AUC \uparrow	58.98 \pm 5.40	55.77 \pm 5.86	64.27 \pm 5.22	51.07 \pm 5.52	59.43 \pm 3.21
SIDER	ROC-AUC \uparrow	55.95 \pm 3.27	57.13 \pm 1.89	56.34 \pm 4.20	57.32 \pm 5.00	53.37 \pm 3.34
HIV	ROC-AUC \uparrow	77.06 \pm 3.16	75.66 \pm 1.26	76.57 \pm 1.39	76.06 \pm 1.06	76.08 \pm 1.33

will interact with others to result in, e.g., different effects on patients. For such tasks, *ConfGen* often leads to significant improvements, providing further evidence for the value of 3D pre-training.

Additionally, for datasets like *bace* with its binding prediction task where 3D information should be valuable, the improvements are only modest. This could suggest that our method does not capture all of the 3D information that is relevant for predicting protein binding, and there is still room for improvement. Another explanation is that a molecule’s geometry is less helpful for *bace* since the geometry of the protein and the binding pocket the molecule has to fit into are not known.

C.2. Different 3D Networks and Ablation

In this section, we justify the design of our 3D network, which we call *Net3D* in this comparison. In Table 11 we compare *Net3D* with different alternative 3D networks, which are the 3D GNNs SMP and EGNN operating on learned node embeddings similar to *Net3D*. Additionally, we ablate the use of our γ function that maps the pairwise distances to a higher dimensional space since it would be an unnecessary complication if it provides no benefit. We call it *Net3D w/o γ* if *Net3D* directly operates on the pairwise distances.

Table 11. Comparison of 3D networks. The MAE of the homo property pre-training and fine-tuning on different halves of QM9. *Net3D w/o γ* refers to dropping the distance encoding of *Net3D*. *Net3D* achieves the best MAE.

METHOD	QM9 MAE
RAND INIT	82.10±0.33
SMP	72.37
EGNN	70.46
<i>Net3D w/o γ</i>	70.34
<i>Net3D</i>	68.96±0.32

In Table 11 we can observe that *Net3D* yields the best downstream performance and that using our γ function is a valuable component of it. Possibly EGNN would benefit similarly from this encoding. SMP’s downstream performance is the worst which could be expected since the 3D input representation which it uses does not uniquely define all the relative positions in a molecule.

We note that SMP is able to distinguish chiral molecules, unlike the other 3D networks, but this advantage cannot be evaluated with our experiments on quantum mechanical properties. Chirality only becomes relevant when considering the interactions between molecules, and in these situations, SMP might be able to leverage its advantage such that our evaluation could be criticized as not entirely fair. Additionally, SMP has much lower memory requirements since it does not suffer from the quadratic complexity of EGNN and *Net3D* in the molecule size. Nevertheless, *Net3D* performs the best, and for drug-like molecules, the quadratic complexity is not problematic.

C.3. Different Methods for Multiple Conformers

We test three main approaches and variations of them for incorporating the 3D information of multiple conformers to justify our choice in the main text. The most straightforward one is **conformer sampling**. We use one of the single conformer setups, but when sampling the batch, we additionally sample $j \in \{1 \dots c_i\}$ and use the single conformer R_i^j . The probability of sampling a conformer is either distributed uniformly (so $1/c_i$ is the probability for each j) or given by the Boltzmann weight of each conformer.

multi3D is the approach from the main text where we include multiple conformers as additional positive pairs in contrastive learning. For each molecule $(G_i, \{R_i^j\}_{j \in \{1 \dots c_i\}})$ we choose the c lowest energy conformers to have a fixed number of them. If there are fewer than c conformers for a molecule ($c_i < c$), then the lowest energy conformer is repeated. For every molecule the 3D network now takes all c conformers $\{R_i^j\}_{j \in \{1 \dots c\}}$ as input and produces their latent 3D representations $\{z_{i,j}^b\}_{j \in \{1 \dots c\}}$ which we can see as additional positive samples. In our contrastive setting, we, therefore, want the similarity between z_i^a and all conformer representations that come from the same molecule $z_{i,j}^b$ to be high. As such, we modify our loss to obtain:

$$\mathcal{L}^{\text{multi3D}} = -\frac{1}{N} \sum_{i=1}^N \left[\log \frac{\sum_{j=1}^c e^{\text{sim}_{\cos}(z_i^a, z_{i,j}^b)/\tau}}{\sum_{\substack{k=1 \\ k \neq i}}^N \sum_{j=1}^c e^{\text{sim}_{\cos}(z_i^a, z_{k,j}^b)/\tau}} \right]. \quad (8)$$

One concern with this formulation is the following. Let us consider a single molecule. The objective of high similarity between the many 3D representations and the single 2D representation might be easier to solve through encoding the same 2D information in the 3D representations instead of capturing the 3D information of all conformers in the single 2D representation. The 2D network would therefore not learn to produce 3D information from its 2D inputs because the mutual information could be maximized through encoded 2D information.

To address this problem, **multi3D+2D** is our third approach. The 2D network is now modified to produce c many latent 2D representations $f^a(G_i) = \{z_{i,j}^a\}_{j \in \{1 \dots c\}}$ which are compared to all 3D representations of the same molecule in a similarity function sim . We simply use this similarity in the loss instead of the cosine similarity. Intuitively, the 2D network now has to produce an embedding for each 3D conformer.

One way to define such a similarity between two same-sized sets of vectors is to use the sum of all pairwise cosine similarities (for brevity we drop the subscript and only write $\{z_{i,j}^a\}$ to mean the set of all representations corresponding to the i -th molecule):

$$sim_{all}(\{z_{i,j}^a\}, \{z_{i,j}^b\}) = \sum_{j=1}^c \sum_{k=1}^c sim_{cos}(z_{i,j}^a, z_{i,k}^b) \quad (9)$$

More principled would be to find the optimal transport matching with the highest cosine similarity, such that one 2D representation always corresponds to one 3D representation. However, this approach was not computationally feasible with the batch sizes we use in contrastive learning. We instead opt for an upper bound on the maximum similarity matching. For every 3D representation, we choose the 2D representation that has the highest similarity. This way, one 2D representation could be associated with multiple 3D embeddings, and we no longer have a mass preserving matching:

$$sim_{max}(\{z_{i,j}^a\}, \{z_{i,j}^b\}) = \sum_{k=1}^c \max_{j \in \{1 \dots c\}} sim_{cos}(z_{i,j}^a, z_{i,k}^b). \quad (10)$$

Beyond these similarity measures, we explore additional ones based on the inverse of different distance functions and asymmetric metrics such as the maximum mean discrepancy (Gretton et al., 2012) or the KL- and JS-Divergence when interpreting the conformer representations as samples from a normal distribution.

Results We evaluate which of the different approaches best leverage the additional conformer’s information to justify our choice for *multi3D* in the main text. Another hypothesis we wish to test is that for smaller molecules such as those in QM9, the ability to make predictions informed by multiple conformers is not as important as for larger drug-like molecules. The reasoning is that a single conformer takes most of the Boltzmann weight for QM9’s molecules due to the fewer degrees of freedom.

We test *conformer sampling*, *multi3D*, and *multi3D+2D* when pre-training on either QMugs or one half of GEOM-Drugs and fine-tuning on QM9 or the other half of GEOM-Drugs. In QMugs we have three diverse conformers available for each molecule which are all used, while for GEOM-Drugs different numbers of conformers are available of which we use the five with the highest Boltzmann weight i.e. lowest energy. If there are fewer than five we duplicate the lowest energy conformer (see Section 4.2 for details). We recall that for the *multi3D+2D* loss sets with as many elements as conformers are produced by the 2D and 3D networks. Both the discussed sim_{all} and sim_{max} are used as similarity measures between those sets. For the *conformer sampling* strategies of using a uniform weighting or sampling conformers according to their Boltzmann weight, we do not evaluate the latter on QMugs since we do not have it available with exactly three conformers per molecule.

In Table 12 we can observe that there are large improvements possible when using multiple conformers. After pre-training on QMugs, the MAE, when predicting the homo property, decreases from 82.57 to 70.77 and from .1966 to .1831 for predicting the Gibbs free energy for GEOM-Drugs. Notably, these improvements are much larger than when pre-training with GEOM-Drugs. This is likely because the GEOM-Drugs dataset contains the lowest energy conformers, and we always use the most probable one with the highest Boltzmann weight when pre-training with a single conformer. Meanwhile, the QMugs dataset contains three diverse conformers per molecule and not the ones with the highest Boltzmann weight. Pre-training with the lowest energy conformer from GEOM-Drugs already captures most of the relevant information, and using more is not as beneficial. However, for QMugs, using the information of all three diverse conformers is crucial.

Similar to the small improvements over the random initialization baseline with GEOM-Drugs, the different methods for using multiple conformers mostly perform the same when pre-training with GEOM-Drugs. When pre-training with QMugs

Table 12. Comparison of strategies for using multiple conformers. The middle double-column shows the results for pre-training on one half of GEOM-Drugs and the right double-column corresponds to pre-training on QMugs, and the second row indicates what dataset was used for fine-tuning. The *Random Init* row shows the performance when training from scratch without any pre-training. For QM9, the reported number is the MAE of the homo property, and for GEOM-Drugs it is the MAE when predicting the ensemble Gibbs free energy. There are large improvements from using multiple conformers, but the differences between the methods are small.

LOSS/ESTIMATOR	GEOM-DRUGS PRE-TRAINING		QMUGS PRE-TRAINING	
	QM9	GEOM-DRUGS	QM9	GEOM-DRUGS
RAND INIT	82.10±0.33	.2035±.0011	82.10±0.33	.2035±.0011
SINGLE CONFORMER	71.66	.1844	82.57	.1966
UNIFORM SAMPLING	70.66	.1823	72.94	.1874
BOLTZMANN SAMPLING	70.93	.1846	X	X
MULTI3D	70.78	.1811	70.77	.1831
<i>sim_{all}</i>	71.11	.1849	72.40	.1936
<i>sim_{max}</i>	70.81	.1896	71.15	.1840

instead, the MAE are overall slightly worse, and we find *multi3D* to perform the best. Note that this is with the slight caveat that the epoch at which pre-training is stopped for all methods was chosen based on where *multi3D* had the lowest MAE.

Due to these results, we consider *multi3D*, and *conformer sampling* with uniform weighting as our best methods since *multi3D* performs slightly better with pre-training on QMugs but *conformer sampling* is simpler and especially uses much less memory. For *multi3D*, all the conformers need to be processed in parallel, and training with more than 5 conformers and a batch size of 500 would not be possible on a 48GB vRAM GPU.

The hypothesis that the downstream performance on the smaller molecules of QM9 would benefit less from using multiple conformers than the molecules of GEOM-Drugs clearly does not hold. Surprisingly, the improvements on the small molecules of QM9 are larger.

C.4. Different Losses

We compare the different losses to estimate and maximize the mutual information. For this purpose, we pre-train PNA on 50 000 molecules from QM9 and another instance on 140 000 molecules of GEOM-Drugs, both with a single conformer. We do so with the Donsker-Varadhan (Hjelm et al., 2019) estimator, the Jensen-Shannon estimator (Hjelm et al., 2019), InfoNCE, and our loss. For our loss, we search over seven temperature parameters $\tau \in [0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.7]$ and choose $\tau = 0.1$.

Table 13. Comparison of mutual information estimators for 3D Infomax. The middle double-column shows the results for pre-training on one half of QM9, and the right double-column corresponds to pre-training on one half of GEOM-Drugs, and the second row indicates what dataset was used for fine-tuning. The *Random Init* row shows the performance when training from scratch without any pre-training. For QM9 the reported number is the MAE of the homo, and for GEOM-Drugs it is the MAE when predicting the ensemble Gibbs free energy.

LOSS/ESTIMATOR	QM9 PRE-TRAINING		GEOM-DRUGS PRE-TRAINING	
	QM9	GEOM-DRUGS	QM9	GEOM-DRUGS
RAND INIT	82.10±0.33	.2035±.0011	82.10±0.33	.2035±.0011
DONSKER-VARADHAN	82.49	.2152	85.46	.2013
JENSEN-SHANNON	80.71	.2078	81.61	.2047
INFONCE	75.81	.1938	79.31	.1894
OUR LOSS	68.96±0.32	.1945	71.66	.1844

In Table 13 we see that 3D pre-training on GEOM-Drugs or QM9 can yield significant improvements for predicting quantum mechanical properties, especially when using InfoNCE and our loss as objectives. These two objectives perform better than the Donsker-Varadhan, and Jensen-Shannon estimator in every case and the Jensen-Shannon objective is superior to the Donsker-Varadhan estimator, which seems to yield no significant improvements over random initialization. The superiority of the Jensen-Shannon loss over the Donsker-Varadhan alternative is in line with the findings of Hjelm et al. (2019) in their different setting on images. While our loss seems to perform better than InfoNCE in three settings, this might be due to the additional investment in searching through temperature parameters for our loss.

C.5. SSL Methods

Here we compare our 3D Infomax pre-training against three additional SSL methods. These are Barlow Twins (Zbontar et al., 2021), multi-modal BYOL (Grill et al., 2020), and VICReg (Bardes et al., 2021). We pre-train these methods on one half of QM9. For a fair comparison, we search through 8 different hyperparameter settings based on the downstream performance on the QM9 homo property. After these method-specific hyperparameters were selected, we tuned every method with a random search over the same search space.

For 3D Infomax, we vary the temperature of our loss τ . When using BYOL we try different decay rates γ for the exponential moving average weight copying. Here, we include $\gamma = 0$ making our setup similar to a multi-modal version of SimSiam (Chen & He, 2020). For Barlow Twins, the hyperparameter is λ weighting the redundancy loss. Lastly, for VICReg we vary μ and ν , the parameters for the variance and the covariance regularization:

1. 3D Infomax with our loss: $\tau \in [0.02, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.7]$ where $\tau = 0.01$ performed the best.
2. Multi-modal BYOL: $\gamma \in [0, 0.0005, 0.001, 0.005, 0.01, 0.03, 0.05, 0.07]$ where $\gamma = 0.005$ performed the best.
3. Barlow-Twins: $\lambda \in [0.002, 0.0039, 0.005, 0.007, 0.01, 0.012, 0.015, 0.02]$ where $\gamma = 0.0039$ performed the best.
4. VICReg: $\lambda = 1$; $\mu \in [1, 0.5]$; $\nu \in [0.02, 0.04, 0.1, 0.3]$ where $\lambda = 1, \mu = 1, \nu = 0.04$ performed the best

Table 14. Comparison of latent space SSL methods. The numbers show the MAE when predicting QM9’s homo property after pre-training on one half of QM9 with the given method and fine-tuning on the other half of QM9. The *Rand Init* column shows the MAE without pre-training and with random weight initialization. 3D Infomax is our best latent space SSL method.

	RANDOM INIT	3D INFOMAX	BYOL	BARLOW TWINS	VICREG
QM9 MAE	82.10±0.33	68.96±0.32	79.16±0.58	82.38±0.48	85.15

The results in Table 14 showcase that 3D Infomax clearly is the superior method in our setting. It decreases the MAE from 82.10 ± 0.33 to 68.96 ± 0.32 while the other methods either lead to no improvement or to the much smaller drop to 79.16 ± 0.58 for BYOL. This is not due to collapse to a constant solution since we can observe a high variance between the representations in a batch for all methods. Furthermore, with the final parameter settings, all methods were able to achieve a low value for their loss during pre-training, both on the training and validation data and there are no optimization issues.

Intuitively, the results can be explained by 3D Infomax being the only method that optimizes a lower bound on the mutual information, which potentially makes it especially fit for our setting. The other approaches have no direct relation to mutual information and instead rely on maximizing a notion of similarity with tricks to prevent collapse. While this might work for conventional SSL, we see no success in our scenario where the rigorous guarantee on maximizing the mutual information seems valuable.

Another reason for the poor performance of BYOL and especially Barlow Twins and VICReg might be that they rely on having symmetric networks to generate the compared representations. In our scenario, we have very little similarity between the architectures with our 2D and 3D networks operating on different modalities. This hypothesis would fit in line with the findings of Bardes et al. (2021) and Zbontar et al. (2021) where introducing asymmetries between the networks hurt performance.

C.6. Pre-training a 3D GNN

We try to use our 3D Infomax setup to pre-train a 3D GNN. For this purpose, we employ SMP (Liu et al., 2021) as 3D network during pre-training with half of the QM9 dataset. We then transfer it’s weights and fine tune them using the accurate 3D conformers of the other half of QM9’s molecules to predict the dataset’s properties. We compare this with SMP trained on the same molecules with randomly initialized weights. The only architectural difference between the networks is that the pre-trained GNN does not use atom features for the reasons explained in the 3D Network paragraph in Section 4.

In Table 15, we find that pre-training improves the 3D GNN’s performance. This may be due to the covalent bonding structure and other 2D edge information that is available during pre-training and which SMP usually cannot use since it employs a radius graph. This is the case even though the pre-trained SMP does not have access to the atom features.

Table 15. MAE for predicting QM9’s molecular properties. SMP is tested with random weight initialization and with the weights obtained from using it as 3D network in our 3D Infomax pre-training setup.

TARGET	SMP RAND INIT	SMP PRE-TRAINED
μ	0.0726	0.0801
α	0.1542	0.1276
HOMO	56.19	44.50
LUMO	43.58	37.48
GAP	85.10	70.45
R2	1.51	1.12
ZPVE	2.69	2.43
c_v	0.0498	0.0421

Pre-training 3D GNNs might be an interesting future direction to attempt beating the state-of-the-art methods for predicting quantum properties with accurate 3D information.

C.7. Cheap Neural Conformers as 3D GNN input

In Section 5.3 we used RDKit’s ETKDG algorithm (Landrum, 2016) to generate inaccurate but cheap and fast conformers and employed them as inputs to the 3D GNN SMP (Liu et al., 2021). Here, we attempt the same with conformers generated by the state-of-the-art deep learning method for conformation generation which is GeoMol (Ganea et al., 2021). For this purpose, we train GeoMol with 50k molecules of QM9 and use it to generate the conformations for the rest of QM9. SMP is then trained on 50k different molecules to predict their properties, either using RDKit’s conformers or those of GeoMol. This enables a fair comparison with 3D Infomax, which uses the same molecules for pre-training that were used to train GeoMol. When visually inspecting some of the conformers generated by GeoMol, we found that they were sometimes of poor quality for molecules with rings and contained outliers with conformations that seem particularly unrealistic.

Table 16. MAE for QM9’s properties. 3D Infomax is tested with three pre-training datasets and compared with the 3D GNN SMP using explicit 3D coordinates. The conformers are generated using the classical method RDKit ETKDG or the learned method GeoMol. Colors indicate improvement (lower MAE) or worse performance compared to the randomly initialized (Rand Init) model.

TARGET	RAND INIT	OUR 3D INFOMAX			RDKIT	GEO MOL
		QM9	DRUGS	QMUGS	SMP	SMP
μ	0.4133±0.003	0.3507	0.3512	0.3668	0.4344	0.6046
α	0.3972±0.014	0.3268	0.2959	0.2807	0.3020	0.8435
HOMO	82.10±0.33	68.96	70.78	70.77	82.51	195.0
LUMO	85.72±1.62	69.51	71.38	78.10	80.36	201.4
GAP	123.08±3.98	101.71	102.59	103.85	114.24	284.1
R2	22.14±0.21	17.39	18.96	18.00	22.63	65.84
ZPVE	15.08±2.83	7.966	9.677	12.06	5.18	17.40
c_v	0.1670±0.004	0.1306	0.1409	0.1208	0.1419	0.5467

Table 16 shows that SMP performs poorly with the conformers generated by GeoMol and using those generated by RDKit is always superior. This is the case even though the average accuracy of GeoMol’s conformers is comparable to that of RDKit ETKDG’s conformers when GeoMol is trained on all of QM9 (Ganea et al., 2021). We hypothesize that the high MAE with GeoMol’s conformers occur since they contained some particularly unrealistic outlier conformations, and SMP is not able to handle those well.

C.8. Batch Size

It is well known that contrastive learning with losses such as NTXent heavily rely on a high number of negative samples - that is large batch sizes. Thus, we evaluate how our loss benefits from large batch sizes in Table 17. We find that our large batch size of 500 is indeed necessary for the good performance of 3D Infomax.

3D Infomax

Table 17. Downstream performance for different pre-training batch sizes. Shown is the MAE when fine-tuning on the homo property of one half of QM9 after pre-training on the other half of QM9.

BATCH SIZE	500	400	300	200	100
QM9 MAE	68.96±0.32	69.81	70.18	71.13	73.65

C.9. Combining Pre-Training Methods

Here we simply use GraphCL’s node drop augmentation for the 2D graph and the 3D information (removing all pairwise distances for a removed atom) with a drop ratio of 0.2 during our 3D pre-training process.

Table 18. Comparison of performance when combining 3D pre-training with conventional pre-training by randomly dropping nodes on the 2D or 3D side (labeled **3D Infomax +**) for various biophysical property OGB datasets. **GraphCL** is another pre-trained baseline. Shown is either the RMSE indicated by ↓ where lower values are better or the ROC-AUC indicated by ↑ where higher values are better. Colors indicate **improvement**, **worse** performance, or **no significant change** compared to the randomly initialized (**Rand Init**) model. 3D Infomax is either on par with random initialization or better. There is no negative transfer as there is with GraphCL.

DATASET	RAND INIT	GRAPHCL	3D INFOMAX	3D INFOMAX +
ESOL↓	0.947±0.038	0.959±0.047	0.894±0.028	0.918±0.037
LIPO↓	0.739±0.009	0.714±0.011	0.695±0.012	0.710±0.007
FRESOLV↓	2.233±0.261	3.744±0.292	2.337±0.227	2.791±0.323
BACE↑	78.13±1.30	77.18±4.01	79.42±1.94	79.28±3.61
BBBP↑	68.27±1.98	71.06±2.00	69.10±1.07	68.64±2.19
TOX21↑	73.88±0.64	78.92±0.61	74.46±0.74	73.73±0.69
TOXCAST↑	63.62±0.48	64.95±0.31	64.41±0.88	63.95±0.38
CLINTOX↑	58.98±5.40	51.07±5.52	59.43±3.21	83.59±3.64
SIDER↑	55.95±3.27	57.32±5.00	53.37±3.34	58.43±1.28
HIV↑	77.06±3.16	76.06±1.06	76.08±1.33	75.38±0.95