Traversing Chemical Space with Active Deep Learning

Derek van Tilborg[1,2] and Francesca Grisoni[1,2*]

[1]Institute for Complex Molecular Systems (ICMS), Department of Biomedical Engineering, Eindhoven University of technology, PO Box 513, 5600 MB Eindhoven, The Netherlands.
[2]Centre for Living Technologies, Alliance TU/e, WUR, UU, UMC Utrecht, Princetonlaan 6, 3584 CB, Utrecht, The Netherlands.
*Corresponding author: f.grisoni@tue.nl.

Abstract

Deep learning is accelerating drug discovery. However, current approaches are often affected by limitations in the available data, *e.g.*, in terms of size or molecular diversity. Active learning is poised to be a solution for drug discovery in low-data regimes. In active learning, a model is updated iteratively by taking multiple smaller screening steps, instead of suggesting many molecules at once with a single model for traditional 'one-shot' screening. This iterative approach aims to improve models during the screening process and can adjust course along the way. Despite pioneering applications, it is currently unclear how active learning holds up to traditional approaches and what the best strategies are for prospective drug discovery applications. In this study, we combine six active learning strategies with two distinct deep learning methods and analyse their performance on three large-scale libraries. For each strategy, we investigate its capacity to discover bioactive hits and explore the chemical space. We found that active learning can achieve up to a six-fold improvement in hit retrieval compared to traditional methods in a low-data scenario. The strategy to acquire new molecules at each cycle resulted to be the primary driver of performance, and it determined the 'molecular journey' throughout screening cycles. Finally, we show that active learning can quickly compensate for a lack of molecular diversity in the starting set. These results advocate for the adoption of active deep learning to accelerate drug discovery in low-data regimes.

Keywords: *Active learning, drug discovery, molecular deep learning, virtual screening*

Introduction

Deep learning is showing increasing promise for drug discovery[1,2]. One of its main applications is virtual screening[3], whereby large commercial libraries (usually consisting of 10^3-10^9 molecular candidates) are prioritized for prospective wet-lab experiments[3–6]. A key bottleneck of deep learning, however, is the need for sufficient data to train a machine learning model (preferably 10^3 molecules and above[7,8]). Unfortunately, available ligand-target interaction data are often limited in size and structural diversity – factors that might hamper the usefulness of deep learning models in practice[8–10]. Furthermore, much of the chemical composition of commercial screening libraries is often highly distinct from the training data[11], resulting in unreliable predictions.

One potential solution to escape the size and diversity limits of the training data is active learning[12–14]. Active learning is based on the principle that a model can achieve greater accuracy with fewer training data if it is "allowed to choose the data from which it learns"[15]. In drug discovery, active learning can be cast into an iterative screening approach (Fig. 1), where instead of performing a single virtual screening experiment, one can test (fewer) molecules across multiple cycles[12–14,16–20]. At each iteration, a selection of molecules is acquired based on model predictions and tested in the wet-lab. The newly obtained experimental data, together with data previously collected, are then used to update the model, aiming to inform the next iteration. By improving models over subsequent iterations, active learning bears promise to identify more bioactive molecules using less resources than standard 'one-shot' virtual screening approaches[12]. Despite this promise and pioneering applications[21–28], the advantages that active learning offers over traditional methods in virtual screening have not been fully quantified yet, especially with a focus on deep learning. This might be why the adoption of active deep learning remains slow in practice. Nonetheless, active learning is expected to gain increasing relevance in the future, *e.g.*, in the context of automated molecule discovery and self-driving labs[29,30].

Although there have been several pioneering active learning studies in the molecular domain[12,21–28,31,32], numerous technical and practical problems remain unaddressed. In this study, we aim to illuminate some of this practical dark matter of active deep learning in drug discovery. We focus on a realistic drug screening scenario and how different components of the active learning pipeline can be tuned to navigate vast screening libraries in search for novel bioactive hits. We compared active deep learning to traditional 'one-shot' screening approaches and similarity-based screening, showing its advantages for chemical space exploration.

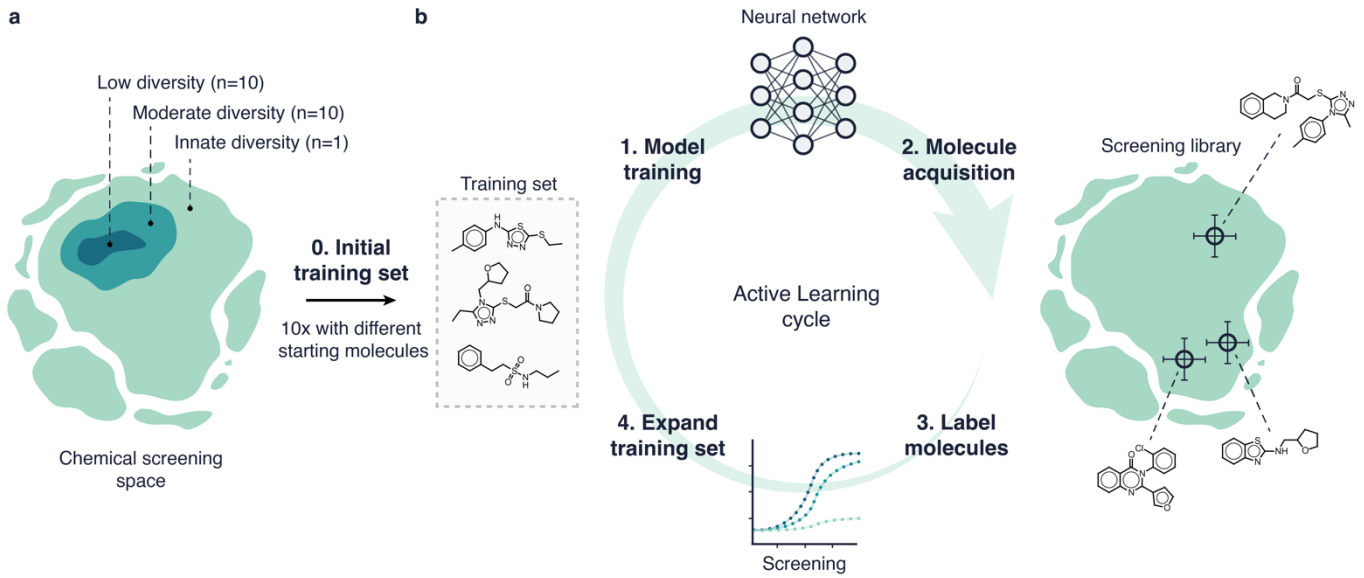

Figure 1: Overview of the experimental setup. a. A starting set of molecules (64) is selected from the screening library (step 0). This starting set can be sampled from the full library (innate structural diversity) or can be sampled from hierarchical subclusters of molecules with moderate or low structural diversity. **b.** The active deep learning cycle. From the training set, a model is trained and used to predict bioactivity on all unlabelled molecules in the screening library (step 1). Using an acquisition function, molecules are selected for follow-up (step 2), and subsequent labelling (step 3). Finally, the labelled molecules are added to the training set for the next cycle (step 4).

Results and Discussion

Active learning setup

In this study, all experiments followed the same setup (Fig. 1), whereby active learning models are used to iteratively query a screening library of 100,000 molecules. Each active learning experiment started with a set of 64 molecules (step 0) sampled from the screening library so that it contained at least one bioactive molecule (*see* Materials and Methods). The active learning screening procedure consisted of four iterative steps:

1. *Training*. A machine learning model is trained on all available training data. The trained model is then used to perform bioactivity predictions on all unlabelled molecules in the screening library.
2. *Acquisition*. From these predictions, 64 molecules are selected for follow-up using a pre-determined strategy ('acquisition function').
3. *Testing*. The acquired molecules are labelled with their corresponding experimental bioactivity (known in advance but not used for model training), to simulate a wet-lab testing procedure.
4. *Update*. All tested molecules are added to the training set, for the next cycle (step 1).

This four-step cycle was repeated until 1,000 molecules were screened. All experiments were performed ten times with different random starting sets. The acquisition size (step 2) was determined through preliminary experiments where we compared acquiring 16, 32, or 64 molecules per cycle and found no significant performance differences (see Supplementary Information).

Using this setup, we investigated the impact of the following factors on the effectiveness of active learning: (a) the structural diversity of the starting set (three levels of diversity), (b) the chosen deep learning approach (two approaches), and (c) the acquisition function used to decide what molecules to select (six acquisition functions). These systematic analyses were carried out on three macromolecular targets, to ensure the robustness and generalizability of the obtained conclusions. These factors are explained below.

Table 1: Summary of the data. Subsets of three datasets from LIT-PCBA[32] were used: Pyruvate kinase M2 (PKM2, agonism), Aldehyde dehydrogenase 1 (ALDH1, inhibition), and Vitamin D receptor (VDR, antagonism).

Dataset	Screening library size	Test set size
PKM2	100,000 (223 hits, 0.2%)	20,000 (44 hits, 0.2%)
ALDH1	100,000 (4986 hits, 5.0%)	20,000 (997 hits, 5.0%)
VDR	100,000 (239 hits, 0.2%)	20,000 (48 hits, 0.2%)

Chosen macromolecular targets

To mimic real drug-screening experiments that rely on large screening libraries, we used three high-throughput screening datasets from LIT-PCBA[33], each for a different biological target. Not only are LIT-PCBA datasets designed to mimic the hit/potency distribution of typical experimental drug screens[33], but they are also large enough to be used as a 'screening library' to simulate prospective active learning campaigns. We selected the three LIT-PCBA datasets containing the most experimentally-validated molecules, which refer to targets of clinical and therapeutic interest[34–37], namely: (a) Pyruvate kinase M2 (PKM2, agonism), (b) Aldehyde dehydrogenase 1 (ALDH1, inhibition), and (c) Vitamin D receptor (VDR, antagonism). For each dataset, 100,000 molecules were randomly extracted, preserving the proportion between active and inactive molecules. These molecules served to construct a screening library, from

which the starting training set and the successive molecule picks are drawn (Table 1 and Supp. Fig. 2). An additional set of 20,000 molecules was randomly selected to serve as an external test set for performance monitoring.

Levels of structural diversity

To examine the effect of structural heterogeneity in the starting set, we artificially created subsets of molecules with different degrees of structural diversity. Diversity was defined through molecular similarity by computing the Tanimoto coefficient on Extended Connectivity Fingerprints[38] (ECFPs), which captures the presence of shared substructures. Using hierarchical clustering, we selected ten clusters in each dataset with moderate structural diversity (average Tanimoto similarity ranging from 0.26 ± 0.03 to 0.28 ± 0.03, Table 2). Within each of these clusters, another subcluster of molecules with low diversity was identified (average similarity higher than 0.36 ± 0.01, Table 2). This gave us a hierarchy of three levels of structural diversity (Fig. 1a): (1) *'innate'* diversity, representing the inherent diversity of the full screening library, (2) *moderate* diversity, and (3) *low* diversity. When constructing the starting set, molecules could, thus, be sampled from areas of each level accordingly. Since our approach is hierarchical, this allowed us to vary molecular diversity in the starting set while staying in the same population of molecules for each of the ten experimental replicates.

Table 2: Hierarchy of molecular diversity in starting sets. Molecular diversity is reported as the mean Tanimoto similarity on Extended Connectivity Fingerprints[37] between all molecules in the starting set. For three datasets, the mean mean similarity of all ten starting sets (n=64) is reported with standard deviations. The higher the mean similarity, the lower the diversity. These values were the best achievable for ten hierarchical subsets large enough to sample the starting set from.

	Similarity ↑		
Diversity ↓	PKM2	ALDH1	VDR
Innate	0.14 ± 0.00	0.13 ± 0.00	0.14 ± 0.00
Moderate	0.27 ± 0.03	0.28 ± 0.04	0.26 ± 0.03
Low	0.46 ± 0.03	0.46 ± 0.02	0.36 ± 0.01

Deep learning models

Two deep learning strategies were used to perform bioactivity predictions. Models were trained using either traditional engineered molecular descriptors or learnable molecular representations[39]:

1. Neural networks (multi-layer perceptron) that learn from molecular fingerprints in the form of Extended Connectivity Fingerprints[38] (ECFPs). These molecular fingerprints encode the presence of radial, atom-centred substructures.
2. Graph neural networks, which learn directly from the molecular graph[4,40]. Molecular graphs are a direct numeric representation of molecular topology, with nodes and edges representing atoms and chemical bonds respectively.

To enable uncertainty estimation, both methods were implemented as approximate Bayesian models through anchored ensembling[41] (see *Materials and Methods*).

Acquisition functions

An important contributor to hit discovery is the so-called acquisition function. The acquisition function determines which molecules are selected for screening and how the training set is expanded; governing hit retrieval and influencing future iterations. Six acquisition functions were investigated as a strategy to select molecules for the next cycle, and for their effect on the active learning performance:

1. *Molecular similarity*. Molecules in the screening library with the highest structural similarity (Tanimoto similarity on ECFPs) to any previously found hit are selected.
2. *Exploitation*. Molecules with the best model predictions are selected (Eq. 6).
3. *Exploration*. Molecules with the most uncertain predictions are selected (Eq. 7), with the goal of 'patching knowledge gaps' in the model.
4. *Mutual Information*[42]. Based on Bayesian Active Learning by Disagreement (BALD)[42], molecules with the lowest mutual information (Eq. 8) are selected from the screening library. Mutual information is low when there are many possible ways of predicting the data with high certainty, given the same model.
5. *Exploitation without retraining*. Molecules are selected with exploitation using a model trained on just the start dataset (Eq. 6). This method resembles traditional 'one-shot' virtual screening where a single model prioritizes the molecules for the whole screening experiment.
6. *Random acquisition*. Molecules are randomly selected from the screening library. This method serves as a control.

Evaluation of active deep learning

The factors investigated in this study were evaluated for their effect on hit enrichment across active learning cycles. This was quantified using the enrichment factor (EF)[43], which captures the ratio between the number of hits found among all acquired molecules and the number of hits expected in selecting the same number of molecules at random from the screening library. EF values larger than 1 indicate methods that can enrich a selection of hits more than a random pick (the higher, the better), while EF values lower than 1 perform worse than random at hit retrieval.

A diverse start

The active learning process is 'seeded' by the starting data that are available on a given target of interest. One has often no control on what molecules are available for training. Here we tested the effect of structural diversity (innate, moderate, and low) in the starting set on the effectiveness of the subsequent active learning cycles. We found that, in general, the molecular diversity of the starting set has little effect on later screening cycles, in terms of diversity of the acquired molecules (Fig. 2) and hit retrieval (Supp. Fig. 3). In fact, the initial structural bias is quickly compensated for in the first 1-5 cycles for most methods (Fig. 2) and converges to the levels of innate

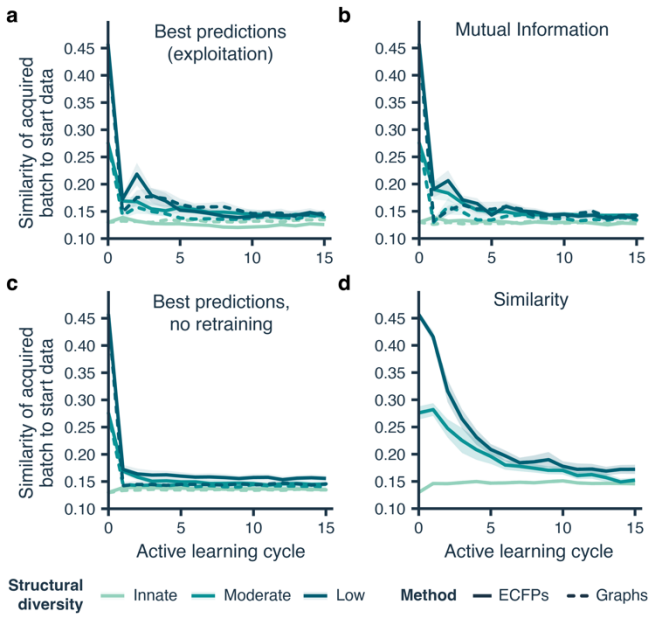

Figure 2: The effect of structural diversity in the starting data. Using the ALDH1 dataset, 64 molecules are acquired each cycle. Shaded areas represent the standard error of 10 experiments with different random starting molecules. Structural diversity is defined as mean Tanimoto similarity within the starting set (Table 2). **a**. Molecule acquisition based on best predictions (exploitation). **b**. Molecule acquisition based on mutual information. **c**. Molecule acquisition based on best predictions (exploitation) without updating models. **d**. Similarity-based molecule acquisition. Since molecules are not acquired based on model predictions, but through molecular similarity, results for ECFP- and graph-based models are identical.

similarity of the screening library (Table 2). Only for similarity-based acquisition, by definition, the structural bias of the starting set lingers. In a few cases (mainly for exploitative and mutual information-based acquisition functions) fewer hits seem to be found when starting with highly similar molecules, although this proved to not be statistically significant in most cases (Supp. Fig. 3). This behaviour is observed regardless of the chosen deep learning approach and datasets. These results indicate that the structural diversity of the starting set does not play a big role in hit identification when screening for several cycles, hence showing the usefulness of iterative train-test-update cycles in navigating chemical libraries effectively.

Choosing an acquisition function

Choosing how to acquire the next iteration of molecules is arguably one of the most crucial factors in determining the hit enrichment. Here, we found that exploitation (selection of the best predicted molecules) and acquisition based on mutual information outperform all other acquisition functions, regardless of the dataset and the deep learning approach (Fig. 3). Active learning achieves a remarkable increase in the number of hits compared to screening the same number of molecules in just one go ('one-shot'). For instance, active learning with exploitation leads to a 2- to 4-fold enrichment compared the corresponding 'one-shot' approach.

Strikingly, iterative similarity-based acquisition (not requiring any machine learning in principle) also yields higher enrichment than 'one-shot' virtual screening across the board, although lower than deep learning. Especially in early iterations, similarity-based acquisition seems to be highly effective, even when compared to the best performing deep learning methods. Finally, exploration (selecting the most uncertain molecules), yields the fewest hits, in line with previous work[12].

Across all datasets and most molecule acquisition functions, ECFP-based models outperform graph-based models in hit enrichment. Graph neural networks, however, achieve a slightly higher prediction accuracy on the PKM2 and ALDH1 datasets (Supp. Table 1 and Supp. Fig. 4). Although graph-based methods tend to find less hits, they exhibit very similar behaviour to the ECFP-based methods in the relative performance of acquisition functions. Therefore, from here on, we will focus solely on ECFP-based models. The results of graph neural networks can be found as Supplementary Information.

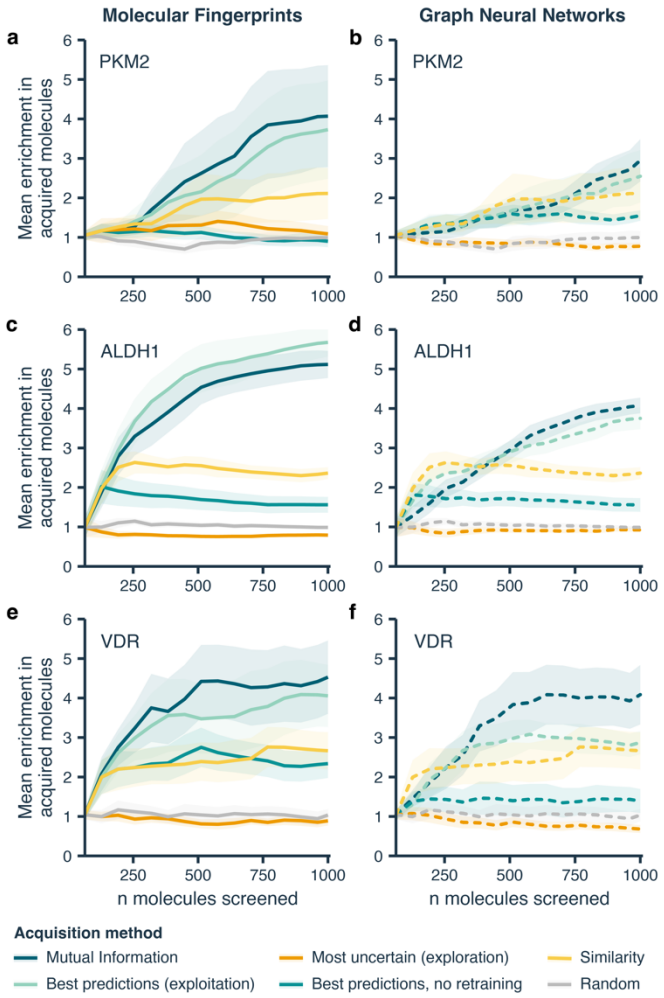

Figure 3: Effect of acquisition functions on hit discovery. Enrichment factor of all acquired molecules with models trained using ECFPs and molecular graphs, reported across three different datasets. Lines represent the average values and shaded areas represent the standard error, computed over 10 experiments with different starting molecules. Molecular fingerprints (ECFPs) are reported on the left-hand column (solid line), and graph neural networks on the right-hand column (dashed line). Every row represents one dataset: **a, b**. PKM2, **c, d**. ALDH1, **e, f**. VDR.

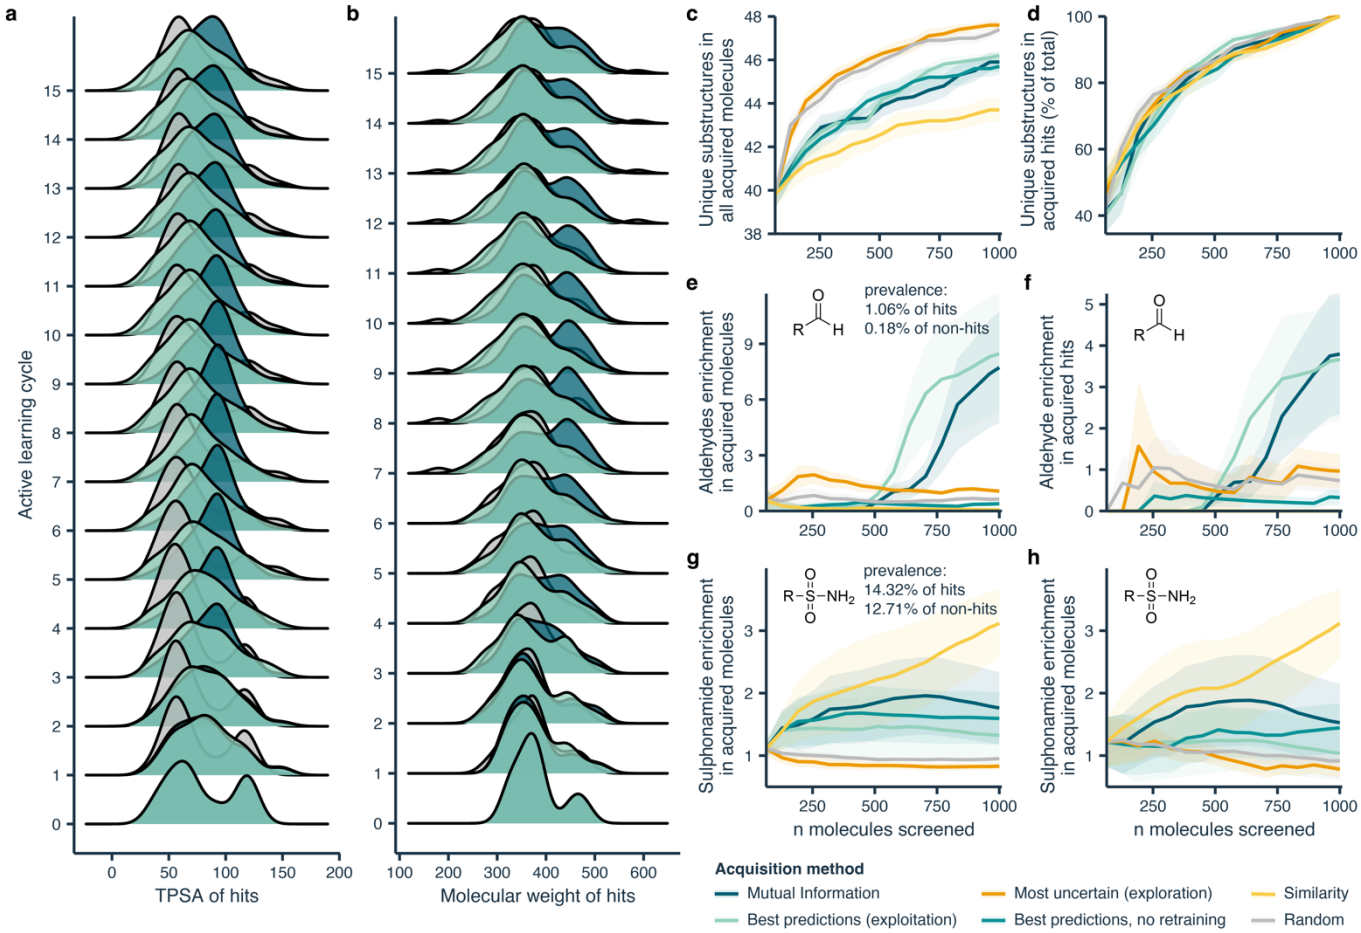

Figure 4: The acquisition function determines the voyage through chemical space. Using the ALDH1 dataset, 64 molecules are acquired each cycle with models trained on ECFPs. **a.** The distribution of total polar surface area (TPSA) of hit molecules throughout active learning cycles. **b**. Molecular weight of hit molecules. **c**. The number of unique chemically relevant molecular substructures found across all acquired molecules. Shaded areas represent the standard error of 10 experiments with different random starting molecules. **d**. The share of unique molecular substructures found in hit molecules per acquisition method. **e**. The enrichment of aldehydes in acquired molecules. **f**. Enrichment of aldehydes among found hits. **g**. The enrichment of sulphonamides in acquired molecules. **h**. Enrichment of sulphonamides among found hits.

Chemical space exploration

Finding structurally novel hits is often one of the objectives of virtual screening since it might increase the chances of success of hit-to-lead optimisation. Hence, we tested the acquisition functions for their ability to find novel hits with (a) different physicochemical properties, and (b) different substructures than the starting training set. We found that active learning is particularly suited to explore novel regions of the chemical space under both perspectives (Fig. 4).

First, different acquisition functions can steer a screening experiment towards populations of molecules with different physicochemical properties than the starting set. An example experiment is shown (Fig. 4a, b), where, starting from the same set of molecules, the physicochemical properties of the acquired hits (*e.g.*, total polar surface area and molecular weight) will progressively move towards different distributions over screening cycles based on the chosen acquisition function. This indicates that different acquisition functions can discover distinct chemistry in the form of their physicochemical properties.

Furthermore, to get more insights into the capacity of different acquisition functions to find structurally novel molecules, we curated a predefined list of 50 chemically-relevant molecular substructures, including functional groups and structural features (Supp. Table 2). This list was used to monitor the discovery of new substructures across active learning cycles. We found that exploration (selection of most uncertain molecules) and random acquisition are the best functions in terms of (global) early enrichment of new substructures (Fig. 4c). Similarity-based acquisition, by definition, is the slowest at finding novel chemistry. When looking at hit molecules only, exploitation and mutual information-based acquisition yield the highest number of new (unique) substructures (Supp. Fig. 5). However, since these acquisition methods find more hits in general, we normalized for the number of hits found with each method to compare the discovery of 'novel substructures per hit' (Fig. 4d). When corrected for their hit discovery efficiency, no single method finds novel substructures faster.

Furthermore, active learning shows an adaptive behaviour in the molecular substructures that are acquired during screening. For example, aldehydes – overrepresented in ALDH1 hits (1.06%) compared to all

molecules (0.18%) – will, at some point, get prioritized by exploitative and mutual information-based acquisition functions (Fig. 4e). In turn, the enrichment of aldehydes found among hits rises accordingly (Fig. 4f). This indicates that relevant structure-activity relationships can be picked up along the way based on newly discovered information. This is not true for methods not based on active learning. In fact, irrelevant substructures might get prioritized purely by their presence in the starting set. For example, in the ALDH1 dataset, similarity-based acquisition starts to prioritize sulphonamides (Fig. 4g), which are common in this dataset, but not overrepresented in hit molecules (occurring in 14.32% of hits and 12.71% of non-hits). The adaptive nature of active learning allows it to change course during screening and might prevent models to get stuck in specific areas of the chemical space.

Conclusion

This systematic study showed that active deep learning achieves a considerable leap in hit-enrichment (up to six-folds) compared to traditional virtual screening, across different real-world scenarios. The biggest factor in hit retrieval is the chosen acquisition function, that controls what molecules are selected for follow-up at each cycle. While graph- and fingerprint-based neural networks show comparable classification performance, the latter show a better capacity to identify novel hits.

The chosen acquisition function also drives the journey through the chemical space of the screening library, in terms of physicochemical properties and substructures. Starting from the same set of molecules, in fact, different acquisition functions will lead to different populations of molecules being sampled. Moreover, the iterative nature of active learning allows it to 'change course' along the way and adapt to newly obtained information when prioritizing molecules based on previous discoveries. This is a clear advantage compared to traditional 'one-shot' approaches.

Finally, we show that the molecular diversity of the initial training set does not seem to affect hit discovery in later active learning cycles. The screening methods we explored in the context of our experimental setup appear to rapidly 'correct' for the lack of structural diversity in the starting set within the next 1-5 screening cycles. This implies that a limited molecular diversity in the training data – which is frequently considered as one of the main obstacles in molecular machine learning – can be overcome by resorting to active deep learning.

Our results corroborate active learning as an effective tool for drug discovery, especially in low-data scenarios. Since the method of acquiring molecules proved to be the main performance driver, there is ample room for improved acquisition functions. Acquisition functions that consider a more chemistry-centred approach, on top of model predictions, bear particular promise to improve molecular novelty among found hits.

We hope that this research will encourage the adoption of active deep learning in prospective studies. We envision that such adoption will accelerate the shift towards fully automated, deep learning-guided, molecule screening. Ultimately, leading to the discovery of better drugs with less resources.

Materials and Methods

Data pre-processing and analysis

Data pre-processing. ALDH1, PKM2 and VDR were downloaded from LIT-PCBA[33] (accessed on August 2023 at https://drugdesign.unistra.fr/LIT-PCBA). SMILES strings were canonicalized using RDkit[44] v. 2022.09.5, after which non-unique SMILES strings and molecules that could not be "kekulized" or featurized (*see Molecule featurization*) were omitted. All molecules were randomly shuffled. Data was randomly split into a screening library (training set) containing 100,000 molecules and test set, containing 20,000 molecules (Table 1).

Molecule featurization. ECFPs were computed with the following settings: length = 1024 bis, radius = 2. For the featurization of molecular graphs, the following atom features were one-hot-encoded: atom type (C, N, O, S, F, Cl, Br, I, P, Si, B, and Se), implicit atom degree, total atom degree (including hydrogens), explicit valence, implicit valence, total valence, implicit hydrogens, total hydrogens, formal charge, and hybridization. Additionally, atom membership of a set of molecular substructures was binary encoded (Supplementary Table 2). ECFPs and all atom features were computed from canonicalized SMILES strings using RDkit[44] v. 2022.09.5.

Clustering. Structurally diverse groups of molecules in the training set were identified using average distance agglomerative clustering, implemented with scikit-learn v. 1.2.1. A distance matrix of all molecules was used based on the 'Tanimoto distance' (computed as 1 – Tanimoto similarity) on ECFPs. The hierarchical clustering dendrogram was cut at a Tanimoto distance (T_d) of T_d = 0.8 to find moderately diverse subclusters and at T_d = 0.61, T_d = 0.70, and T_d = 0.70 for PKM2, VDR, and ALDH1 respectively to find smaller, low diversity, sub-subclusters. These specific cut-off values were chosen so that each dataset contained 10 clusters. All subclusters had a minimal size of 128 molecules and contained a single sub-subcluster of 64 molecules. Additionally, clusters should contain a minimal number of hits, comparable to the proportion of hits in the full dataset, as determined by:

$$n_{hits} > n_{min}\mathbb{E}[Y_{screen}], \quad (1)$$

where n_{hits} is the number of hits in the subcluster, n_{min} is the minimum size of the subcluster, and Y_{screen} is the binary class vector of the full screening library.

Molecular substructure assignation. A set of 50 unique SMILES Arbitrary Target Specification (SMARTS) patterns were used to identify substructures (Supplementary Table 2) using RDkit[44] v. 2022.09.5.

Active learning
For each active learning experiment, 64 molecules were selected from the screening library. These molecules are then removed from the library. Every set starts with a random hit molecule, after which the remaining molecules are randomly sampled using a uniform distribution. Until the budget of 1000 molecules was depleted, a model M was trained on the training set, after which bioactivity predictions were made on the screening library. Using one of five acquisition methods (see Acquisition methods), n molecules are labelled and moved to the training set.

Deep learning
Neural network implementation and architectures. All models were implemented using PyTorch[45] v. 1.12.1 and PyTorch Geometric[46] v. 2.2.0. Graph Convolutional Networks (GCN) consisted of an atom embedding layer, followed by three layers of graph convolutions[47] with batch normalization[48]. Global pooling by summing was then used to get molecular embeddings from atom embeddings. Finally, a multi-layer perceptron (MLP) was used with three fully connected layers with batch normalization[48]. The same MLP architecture was used for models trained on ECFPs. All models used a hidden size of 1024 neurons and were optimized for 50 epochs using the Adam algorithm, with a learning rate of 3 x 10^{-4} with mixed precision and a minibatch size of 64. During training, minibatches were resampled based on their class with

$$P_c = 1 - \frac{n_c}{N}, \quad (2)$$

where P_c is the sampling probability for a class c, n_c is the number of samples of class c, and N is the total number of samples.

Uncertainty estimation. For uncertainty estimation, anchored ensembling was used[41]. This ensembling method produces predictive posterior distributions that closely approximate exact Bayesian methods. We used an ensemble of $M = 10$ models. For each model, $m \in \{1 \dots M\}$, its parameters θ_m are regularized with a set of anchored parameters $\theta_{anchor,m}$. Each model is initiated with distinct θ_{anchor}, controlled by random seeding. The classification loss in our implementation is defined as:

$$\mathcal{L}oss = -\frac{1}{N}\sum_{i=1}^{N} \log\left(p_m(y_i|x_i)\right) + \frac{\lambda}{N}\left\|\theta_m - \theta_{anchor,m}\right\|^2, \quad (3)$$

where λ is a regularization coefficient (set to 3×10^{-4}). For estimating the expected value $\mathbb{E}$ of a molecule x_i, we take the mean prediction across all models in the ensemble,

$$\mathbb{E}(y_i|x_i) = \frac{1}{M}\sum_{m=1}^{M} p_m(y_i|x_i). \quad (4)$$

Similarly, the prediction uncertainty for a molecule x_i is defined as the mean entropy $\mathbb{H}$ over the ensemble:

$$\mathbb{H}(y_i|x_i) = -\frac{1}{M}\sum_{m=1}^{M} p_m(y_i|x_i)\log(p_m(y_i|x_i)). \quad (5)$$

Acquisition functions
Five acquisition functions were used to select follow-up molecules at each iteration. Each acquisition function (a) greedily selects n molecules based on model prediction on the screening library.

1. *Similarity-based*: samples are selected based on their highest Tanimoto similarity (computed with ECFPs; with 1024 bits and a radius of 2) to any previously acquired hit compound).

2. *Exploitative*: the best predicted samples are selected with:

$$a_{exploit} = \text{argmax}_n\left(\mathbb{E}(y|x)\right). \quad (6)$$

3. *Explorative*: most uncertain samples are selected with:

$$a_{explore} = \text{argmax}_n\left(\mathbb{H}(y|x)\right). \quad (7)$$

4. *Mutual Information*: selects samples with low mutual information ($\mathbb{I}$) with:

$$a_{\mathbb{I}} = \text{argmin}_n(\mathbb{H}(y|x) - \mathbb{E}_M[\mathbb{H}(y|x,\theta)]). \quad (8)$$

Based on Bayesian Active Learning by Disagreement (BALD)[42], the left term represents the entropy of the model predictions (uncertainty) and the right term represents the expected value of the entropy of the model predictions for each draw of the model parameters (*i.e.*, 'disagreement' between the different models in the ensemble)[49]. To have low mutual information, the model must have many ways of explaining the data with high certainty, *i.e.*, low uncertainty and high disagreement.

5. *Random*: samples are selected from a uniform probability distribution.

Acknowledgements

This research was co-funded by the European Union (ERC, ReMINDER, 101077879). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. The authors also acknowledge support from the Irene Curie Fellowship, the Centre for Living Technologies, and SURF (NWO grant EINF-5379). The authors would like to thank A. Ortiz-Pérez for proofreading the manuscript, L. Rossen for early discussion on active learning, and R. Özçelik for his suggestion of agglomerative clustering.

Data and code availability

The datasets[33] and the Python code to replicate and extend our study, is freely available on GitHub at the following URL: https://github.com/molML/traversing_chem_space.

Author Contributions

Conceptualization: D.v.T and F.G. Data curation: D.v.T. Formal analysis: D.v.T. Methodology: D.v.T and F.G. Software: D.v.T. Visualisation: D.v.T. Writing – original draft: D.v.T. Writing – review and editing: D.v.T and F.G. All authors have given approval to the final version of the manuscript.

References

1. Qureshi, R. *et al.* AI in drug discovery and its clinical relevance. *Heliyon* **9**, e17575 (2023).
2. Jiménez-Luna, J., Grisoni, F., Weskamp, N. & Schneider, G. Artificial intelligence in drug discovery: recent advances and future perspectives. *Expert Opinion on Drug Discovery* **16**, 949–959 (2021).
3. Walters, W. P., Stahl, M. T. & Murcko, M. A. Virtual screening—an overview. *Drug Discovery Today* **3**, 160–178 (1998).
4. Stokes, J. M. *et al.* A Deep Learning Approach to Antibiotic Discovery. *Cell* **180**, 688-702.e13 (2020).
5. Liu, G. *et al.* Deep learning-guided discovery of an antibiotic targeting Acinetobacter baumannii. *Nat Chem Biol* (2023).
6. Neves, B. J. *et al.* Discovery of New Anti-Schistosomal Hits by Integration of QSAR-Based Virtual Screening and High Content Screening. *J. Med. Chem.* **59**, 7075–7088 (2016).
7. Deng, J. *et al.* A systematic study of key elements underlying molecular property prediction. *Nat Commun* **14**, 6395 (2023).
8. Van Tilborg, D., Alenicheva, A. & Grisoni, F. Exposing the Limitations of Molecular Machine Learning with Activity Cliffs. *J. Chem. Inf. Model.* **62**, 5938–5951 (2022).
9. Volkov, M. *et al.* On the Frustration to Predict Binding Affinities from Protein–Ligand Structures with Deep Neural Networks. *J. Med. Chem.* **65**, 7946–7958 (2022).
10. Wallach, I. & Heifets, A. Most Ligand-Based Classification Benchmarks Reward Memorization Rather than Generalization. *J. Chem. Inf. Model.* **58**, 916–932 (2018).
11. Zabolotna, Y. *et al.* Chemspace Atlas: Multiscale Chemography of Ultralarge Libraries for Drug Discovery. *J. Chem. Inf. Model.* **62**, 4537–4548 (2022).
12. Reker, D. & Schneider, G. Active-learning strategies in computer-assisted drug discovery. *Drug Discovery Today* **20**, 458–465 (2015).
13. Graff, D. E., Shakhnovich, E. I. & Coley, C. W. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chem. Sci.* **12**, 7866–7881 (2021).
14. Reker, D. Practical considerations for active machine learning in drug discovery. *Drug Discovery Today: Technologies* **32–33**, 73–79 (2019).
15. Settles, B. *Active Learning Literature Survey*. https://minds.wisconsin.edu/handle/1793/60660 (2009).
16. Murphy, R. F. An active role for machine learning in drug development. *Nat Chem Biol* **7**, 327–330 (2011).
17. Fujiwara, Y. *et al.* Virtual Screening System for Finding Structurally Diverse Hits by Active Learning. *J. Chem. Inf. Model.* **48**, 930–940 (2008).
18. Pyzer-Knapp, E. O. Bayesian optimization for accelerated drug discovery. *IBM J. Res. & Dev.* **62**, 2:1-2:7 (2018).
19. Gentile, F. *et al.* Deep Docking: A Deep Learning Platform for Augmentation of Structure Based Drug Discovery. *ACS Cent. Sci.* **6**, 939–949 (2020).
20. Gusev, F., Gutkin, E., Kurnikova, M. G. & Isayev, O. Active Learning Guided Drug Design Lead Optimization Based on Relative Binding Free Energy Modeling. *J. Chem. Inf. Model.* **63**, 583–594 (2023).
21. Desai, B. *et al.* Rapid Discovery of a Novel Series of Abl Kinase Inhibitors by Application of an Integrated Microfluidic Synthesis and Screening Platform. *J. Med. Chem.* **56**, 3033–3047 (2013).
22. Xue, D. *et al.* Accelerated search for materials with targeted properties by adaptive design. *Nat Commun* **7**, 11241 (2016).
23. Reker, D., Schneider, P. & Schneider, G. Multi-objective active machine learning rapidly improves structure–activity models and reveals new protein–protein interaction inhibitors. *Chem. Sci.* **7**, 3919–3927 (2016).
24. Yuan, R. *et al.* Accelerated Discovery of Large Electrostrains in $BaTiO_3$-Based Piezoelectrics Using Active Learning. *Advanced Materials* **30**, 1702884 (2018).
25. Mekki-Berrada, F. *et al.* Two-step machine learning enables optimized nanoparticle synthesis. *npj Comput Mater* **7**, 55 (2021).
26. Ortiz-Perez, A., Van Tilborg, D., Van Der Meel, R., Grisoni, F. & Albertazzi, L. Machine learning-guided high throughput nanoparticle design. Preprint at https://doi.org/10.26434/chemrxiv-2023-sqb5c (2023).
27. Besnard, J. *et al.* Automated design of ligands to polypharmacological profiles. *Nature* **492**, 215–220 (2012).

28. Borkowski, O. *et al.* Large scale active-learning-guided exploration for in vitro protein production optimization. *Nat Commun* **11**, 1872 (2020).
29. Abolhasani, M. & Kumacheva, E. The rise of self-driving labs in chemical and materials sciences. *Nat. Synth* **2**, 483–492 (2023).
30. Seifrid, M. *et al.* Autonomous Chemical Experiments: Challenges and Perspectives on Establishing a Self-Driving Lab. *Acc. Chem. Res.* **55**, 2454–2466 (2022).
31. Graff, D. E. *et al.* Self-Focusing Virtual Screening with Active Design Space Pruning. *J. Chem. Inf. Model.* **62**, 3854–3862 (2022).
32. Bellamy, H., Rehim, A. A., Orhobor, O. I. & King, R. Batched Bayesian Optimization for Drug Design in Noisy Environments. *J. Chem. Inf. Model.* **62**, 3970–3981 (2022).
33. Tran-Nguyen, V.-K., Jacquemard, C. & Rognan, D. LIT-PCBA: An Unbiased Data Set for Machine Learning and Virtual Screening. *J. Chem. Inf. Model.* **60**, 4263–4273 (2020).
34. Jiang, F. *et al.* Aldehyde Dehydrogenase 1 Is a Tumor Stem Cell-Associated Marker in Lung Cancer. *Molecular Cancer Research* **7**, 330–338 (2009).
35. Yang, C. *et al.* Aldehyde dehydrogenase 1 (ALDH1) isoform expression and potential clinical implications in hepatocellular carcinoma. *PLoS ONE* **12**, e0182208 (2017).
36. Palsson-McDermott, E. M. *et al.* Pyruvate Kinase M2 Is Required for the Expression of the Immune Checkpoint PD-L1 in Immune Cells and Tumors. *Front. Immunol.* **8**, 1300 (2017).
37. Plum, L. A. & DeLuca, H. F. Vitamin D, disease and therapeutic opportunities. *Nat Rev Drug Discov* **9**, 941–955 (2010).
38. Rogers, D. & Hahn, M. Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.* **50**, 742–754 (2010).
39. Zhang, Y. & Lee, A. A. Bayesian semi-supervised learning for uncertainty-calibrated prediction of molecular properties and active learning. *Chem. Sci.* **10**, 8154–8163 (2019).
40. Kearnes, S., McCloskey, K., Berndl, M., Pande, V. & Riley, P. Molecular graph convolutions: moving beyond fingerprints. *J Comput Aided Mol Des* **30**, 595–608 (2016).
41. Pearce, T., Leibfried, F. & Brintrup, A. Uncertainty in Neural Networks: Approximately Bayesian Ensembling. in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics* 234–244 (PMLR, 2020).
42. Houlsby, N., Huszár, F., Ghahramani, Z. & Lengyel, M. Bayesian Active Learning for Classification and Preference Learning. Preprint at http://arxiv.org/abs/1112.5745 (2011).
43. Truchon, J.-F. & Bayly, C. I. Evaluating Virtual Screening Methods: Good and Bad Metrics for the "Early Recognition" Problem. *J. Chem. Inf. Model.* **47**, 488–508 (2007).
44. Landrum, G. RDKit: Open-source cheminformatics. (2006).
45. Paszke, A. *et al.* PyTorch: An Imperative Style, High-Performance Deep Learning Library. in *Advances in Neural Information Processing Systems* (eds. Wallach, H. et al.) vol. 32 (Curran Associates, Inc., 2019).
46. Fey, M. & Lenssen, J. E. Fast Graph Representation Learning with PyTorch Geometric. Preprint at http://arxiv.org/abs/1903.02428 (2019).
47. Kipf, T. N. & Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. Preprint at http://arxiv.org/abs/1609.02907 (2017).
48. Ioffe, S. & Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. Preprint at http://arxiv.org/abs/1502.03167 (2015).
49. Kirsch, A., van Amersfoort, J. & Gal, Y. BatchBALD: Efficient and Diverse Batch Acquisition for Deep Bayesian Active Learning. in *Advances in Neural Information Processing Systems* (eds. Wallach, H. et al.) vol. 32 (Curran Associates, Inc., 2019).
50. McInnes, L., Healy, J. & Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. Preprint at https://doi.org/10.48550/arXiv.1802.03426 (2020).
51. Citovsky, G. *et al.* Batch Active Learning at Scale. in *Advances in Neural Information Processing Systems* (eds. Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P. S. & Vaughan, J. W.) vol. 34 11933–11944 (Curran Associates, Inc., 2021).