#### 000 PHARMACOPHORE-BASED DESIGN BY LEARNING ON 001 002 VOXEL GRIDS

Anonymous authors

Paper under double-blind review

#### ABSTRACT

Ligand-based drug discovery (LBDD) relies on making use of known binders to a protein target to find structurally diverse molecules similarly likely to bind. This process typically involves a brute force search of the known binder (query) against a molecular library using some metric of molecular similarity. One popular approach overlays the pharmacophore-shape profile of the known binder to 3D conformations enumerated for each of the library molecules, computes overlaps, and picks a set of di-016 verse library molecules with high overlaps. While this virtual screening workflow has had considerable success in hit diversification, scaffold hopping, and patent busting, it scales poorly with library sizes and restricts candidate generation to existing library compounds. Leveraging recent advances in voxel-based generative modelling, we propose a pharmacophore-based generative model and workflows that address the scaling and fecundity issues of conventional pharmacophore-based virtual screening. We introduce *VoxCap*, a voxel captioning method for generating SMILES strings from voxelised molecular representations. We propose two workflows as practical use cases as well as benchmarks for pharmacophore-based generation: de-novo design, in which we aim to generate new molecules with high pharmacophore-shape similarities to query molecules, and fast search, which aims to combine generative design with a cheap 2D substructure similarity search for efficient hit identification. Our results show that VoxCap significantly outperforms previous methods in generating diverse de-novo hits. When combined with our fast search workflow, VoxCap reduces computational time by orders of magnitude while returning hits for all query molecules, enabling the search of large libraries that are intractable to search by brute force.

031 032 033

034

003 004

006

008 009

010 011

012

013

014

015

017

018

019

021

025

026

027

028

029

#### INTRODUCTION 1

Ligand-based drug discovery (LBDD) aims to identify and develop pharmacologically active compounds by making use of molecules that are known to interact with a given biological target (Kurogi and Guner, 2001). LBDD-based methods can be useful in the absence of the target's structure 037 (Mason et al., 2001; González et al., 2009; Bernard et al., 2005; Loew et al., 1993), and can be complementary to structure-based drug design when target information is present, for tasks such as virtual screening and lead diversification (Grebner et al., 2020). Acharya et al. (2011) presents a 040 detailed review of LBDD. Central to LBDD is the concept of a pharmacophore, which is defined as a 041 feature of a molecule, located in 3D space, that may be involved in binding (Pearlman, 1993; McNaught 042 et al., 1997; Chang and Swaan, 2006). Examples include hydrogen bond donors and acceptors. 043

Drug discovery campaigns typically rely on pharmacophore-shape based design for finding a diverse 044 pool of molecules that are likely to bind to a target in the same way as a query molecule that has already been shown to bind. Having a diverse pool of binders increases the chances of finding a molecule 046 that has minimal liabilities and has the potential to be optimised into a successful drug (Hawkins 047 et al., 2007). This is conventionally done by running a brute force search against an existing library of 048 molecules to find molecules whose 3D shape and arrangement of pharmacophore features in 3D space (i.e. pharmacophore-shape profiles) have high overlap in 3D space with the pharmacophore-shape profile of the query molecule (Rush et al., 2005); such molecules are known as hits with respect to 051 the query molecule. The process involves three distinct steps: (i) generating conformations for the query molecule and all molecules in the library, (ii) computing the pharmacophore-shape profiles of 052 these molecules, and (iii) aligning these profiles in 3D space and computing the overlap. This workflow (particularly the steps post-conformer generation/retrieval) is typically done using software such as

# OpenEye's Rapid Overlay of Chemical Structures (ROCS) (Rush et al., 2005; OpenEye, Cadence Molecular Sciences, 2024).

There are two main issues with the prevailing pharmacophore-based virtual screening workflows. First, 057 the process of aligning the query and library molecules and computing their overlap is time-consuming and expensive. This is particularly true because this process has to be carried out for each new query molecule in a project, of which we would ideally like there to be hundreds in the absence of major compu-060 tational constraints. This bottleneck limits the number of query molecules that are screened in a typical 061 drug discovery project, limiting the chances of finding a binding candidate molecule with promising 062 properties and minimal liabilities. Second, candidate generation for any query is limited to the same set 063 of compounds that is the existing library. The practical advantage of library screening is that molecules 064 in these libraries are usually far more synthetically accessible (either already in-stock in commercial or corporate collections, or easily synthesizable via established routes and readily accessible building 065 blocks) than de-novo generated molecules. However, this constraint does prevent the discovery of new 066 compounds that may have better properties for the given task than those already present in the library. 067

In this work, we propose a generative model and workflows to address these issues. We make use
of recent advances in voxel-based generative modelling, which have demonstrated the success of
processing and generating molecules in voxel space (Pinheiro et al., 2023; 2024). Concretely, we
propose *VoxCap*, a voxel-based generative model for pharmacophore/ligand-based drug discovery.
VoxCap is a voxel captioning method that generates SMILES strings (Weininger, 1988) from an input
voxelised representation of a molecule. We use VoxCap to generate SMILES strings corresponding
to an input voxelised pharmacophore-shape profile.

We propose and evaluate two workflows: *de-novo* design, where we generate new molecules that are not in the training or test sets, with high pharmacophore-shape similarities to query molecules from the test set; and *fast search*, where we use *de-novo* design in conjunction with a fast 2D substructure similarity search to quickly and efficiently find hits in 3D space from a given library. Our evaluation shows that VoxCap outperforms previous work, by up to an order of magnitude, on the number of diverse *de-novo* hits generated given a query molecule. We also find that VoxCap, in conjunction with our *fast search* workflow, reduces the computational time/cost of inference by orders of magnitude while returning an appreciable number of hits for each query molecule.

082

#### 2 RELATED WORK

084 085

087

## 2.1 IMAGE CAPTIONING

880 Deep learning has significantly automated the generation of captions from natural images. Vinyals et al. (2015) proposed the first end-to-end deep learning model to generate a caption from an input image. 089 They use a convolutional neural network (CNN) (LeCun et al., 1998) to extract features from an input 090 image, and a long short term memory network (LSTM) (Hochreiter and Schmidhuber, 1997) to decode 091 this representation into a caption. Prior to this, Mao et al. (2014) and Kiros et al. (2014) proposed 092 models that separately encoded images and text to rank or generate new captions. Xu et al. (2016) 093 built on Vinyals et al. (2015)'s model by introducing an attention mechanism into the model to attend to 094 different aspects of the input image. Continuing in the path of moving to attention-based architectures, Li et al. (2019) proposed a captioning model based on the Transformer (Vaswani et al., 2017). VoxCap 096 modifies such approaches to operate on 3D inputs, thereby producing strings from voxel grids.

097 098

099

#### 2.2 VOXEL-BASED MOLECULE GENERATION

100 While point-cloud based molecular generation remains widely used, recent work has focused on 101 voxel-based molecular generative models. Ragoza et al. (2020) first proposed a voxel-based generative 102 model for molecules using a variational autoencoder (VAE) (Kingma and Welling, 2022). Skalic et al. 103 (2019) proposed a VAE-LSTM based model to generate SMILES strings given an input voxelised 104 atomic representation of a molecule. Pinheiro et al. (2023) proposed a voxel-based generative model 105 based on the neural empirical Bayes framework (Saremi and Hyvärinen, 2019), implemented using a UNet-based (Ronneberger et al., 2015) architecture. (Pinheiro et al., 2024) extended this framework 106 to the conditional setting for binder generation, generating voxelised molecules conditioned on 107 the voxelised representation of a protein pocket. Ragoza et al. (2022) also proposed a binding-site

conditioned voxel-based generative model using a conditional VAE-based architecture (Kingma and Welling, 2022; Sohn et al., 2015).

110 111

112

124

137

147

155

156

157 158 2.3 PHARMACOPHORE-BASED MOLECULAR DESIGN

113 Typical small molecule pharmacophore-based virtual screening workflows involve searching existing 114 databases for molecules that match the pharmacophore and shape profiles of various lead compounds. Rush et al. (2005) proposed software to define, align and compute the overlap of pharmacophore-shape 115 profiles, as well as a workflow to apply this to LBDD. Grebner et al. (2020) studied the relationship 116 between database size and matching molecules found. Skalic et al. (2019) generated SMILES strings 117 from voxelised atomic density input and conditioning pharmacophore information, and evaluated 118 these against reference sets using distributional metrics such as uniqueness, Fréchet ChemNet Distance 119 (Preuer et al., 2018) and distributional scaffold frequency similarity. Zhu et al. (2023) proposed a 120 conditional VAE-based model to generate SMILES strings conditioned on stochastically generated 121 pharmacophore graphs. They carry out evaluation by measuring a graph-matching score between query 122 and generated molecules, and measuring properties such as docking score of generated molecules 123 against the pocket from which the hit molecule was obtained.

<sup>125</sup> 126 3 METHOD

127 VoxCap generates SMILES strings from an input voxelised representation of a molecule. Our task here 128 is to generate SMILES strings from a molecule's pharmacophore and shape profile. Our 3D voxelised 129 molecular representation builds on previous works on voxelised molecule generation (Pinheiro et al., 130 2023; 2024). We represent a molecule with a Gaussian-like density centered around the location 131 of each of the pharmacophores of the molecule, with each pharmacophore type having a separate 132 channel. We use 6 pharmacophore types: hydrogen bond donor, hydrogen bond acceptor, cation, 133 anion, aromatic ring and hydrophobe. We also use an additional shape channel, in which Gaussian densities are centered at the 3D positions of each of the atoms in the molecule, but in one channel 134 and with all densities having the same radius. An example voxelised pharmacophore-shape profile 135 is given in Figure 1. Section B of the Appendix gives further details on the voxelisation. 136

#### 138 3.1 TRAINING OBJECTIVE

Given a molecule *m* including its conformation, from a dataset  $M = \{m_n\}_{n=1}^N$ , where *N* is the number of molecules in the dataset, we compute its tokenised SMILES string sequence  $\mathbf{u} = (u^{(1)}, \dots, u^{(T)})$ , where *T* is the length of the sequence and each element of **u** is an integer  $r \in \mathbb{Z} : 1 \le r \le R$ representing a token from the vocabulary  $\mathcal{V} = \{v_r\}_{r=1}^R$  of size *R*. Each vocabulary element  $v_r$  is a sequence of characters; we use the vocabulary and tokeniser provided in the DeepChem package (Ramsundar et al., 2019). We also compute the molecule's voxelised pharmacophore-shape profile  $\xi = f_{\xi}(m) \in \mathbb{R}^{C \times d \times d \times d}$ , with number of channels C = 7 and each side of the voxel grid having dimensionality *d*, using a deterministic non-differentiable function  $f_{\xi}$ .

Our objective during training is to predict the ground truth SMILES string sequence u from the input voxelised representation  $\xi$ . The task is to maximise the likelihood of the data with respect to the parameters  $\theta$  of our model. Our model is autoregressive and factorises the likelihood of each datapoint as follows:

$$p(\mathbf{u}|\xi;\theta) = \prod_{t=1}^{T} p(u^{(t+1)}|\mathbf{u}^{(1:t)},\xi;\theta).$$
(1)

The loss  $\mathcal{L}$  for each datapoint is thus given by:

 $\mathcal{L}(\mathbf{u}|\boldsymbol{\xi};\boldsymbol{\theta}) = -\log p(\mathbf{u}|\boldsymbol{\xi};\boldsymbol{\theta}) = -\sum_{t=1}^{T} \log p(u^{(t+1)}|\mathbf{u}^{(1:t)},\boldsymbol{\xi};\boldsymbol{\theta}).$ (2)

As implied by this equation, the model is trained using teacher forcing. At each time step t the model uses the softmax function to output a distribution  $p(u^{(t+1)} = v_r | \mathbf{u}^{(1:t)}, \xi; \theta) = \frac{\exp(f(\mathbf{u}^{(1:t)}, \xi; \theta)_r)}{\sum_{j=1}^R \exp(f(\mathbf{u}^{(1:t)}, \xi; \theta)_j)}$ where  $f(.; \theta) : (\mathbb{R}^* \times \mathbb{R}^{C \times d \times d \times d}) \to \mathbb{R}^R$  is a neural network with parameters  $\theta$  mapping a sequence



also pass it in to this LSTM at every time step, concatenated with the embedded token for that timestep.
 A diagram of the model is given in Supplementary Figure 1.

219 220

221

222

4 **EXPERIMENTS** 

4.1 DE-NOVO GENERATION FROM PHARMACOPHORE-SHAPE INPUT

We carry out experiments to evaluate the ability of our model to generate a diverse set of molecules matching an input pharmacophore-shape profile for the purpose of ligand-based drug discovery. To evaluate the overlap between the pharmacophore-shape profiles of two molecules, we compute OpenEye's ROCS Tanimoto Combo score. For the purposes of this paper, we abstract this as a function TC( $\cdot$ , $\cdot$ ) :  $\mathcal{M} \times \mathcal{M} \rightarrow [0,2]$ , where  $\mathcal{M}$  is the space of all possible molecules with their conformations. TC hence takes in two molecules and outputs a score between 0 and 2, with 0 representing no overlap and 2 representing perfect overlap.

We carry out the evaluation on two datasets: *GEOM-drugs* (Axelrod and Gomez-Bombarelli, 2022) and the subset of *ChEMBL* (Mendez et al., 2019) defined in the GuacaMol benchmark (Brown et al., 2019), commonly used in the machine learning literature.

GEOM-drugs consists of approximately 430,000 drug-like molecules with up to 181 heavy atoms of type { C, O, N, F, S, Cl, Br, P, I and B }. The median number of heavy atoms per molecule is 45 and more than 99% of molecules have less than 80 heavy atoms. The dataset consists of multiple conformations for each molecule. We carry out experiments using different resolutions and grid sizes, and for our final model choose a grid of dimension 48<sup>3</sup> with a resolution of 0.35 A. This volume covers over 99.8% of all points in the dataset. We perform the same preprocessing and dataset split as (Vignac et al., 2023), resulting in a 1.1M/146K/146K train/validation/test split.

The ChEMBL (subset) dataset consists of approximately 1,600,000 drug-like molecules with up to
88 heavy atoms of type { C, O, N, F, S, Cl, Br, P, I, B, Se and Si }. The median number of heavy atoms
per molecule is 27. As the dataset is given as SMILES strings, we use RdKit (Landrum et al., 2023)
to generate a 3D conformation for each molecule. We do this by using ETKDG (Riniker and Landrum,
2015) to get initial coordinates for the atoms followed by optimisation with reference to the Merck
Molecular Force Field (Halgren, 1996). Our final model uses a grid of dimension 64<sup>3</sup> with resolution
0.35. We use the same 1.3M/80k/240k train/validation/test split as Brown et al. (2019).

For each dataset, we carry out evaluation on 100 randomly selected 'query' molecules from the test set. 248 A conformation is associated with each of these query molecules. As a baseline, for each of these query 249 molecules we randomly sample 1000 molecules from the test set. We obtain multiple conformations 250 for each of these sampled molecules using OpenEye, overlay them with the candidate molecule and 251 compute the Tanimoto Combo score. For each sampled molecule, we retain the highest Tanimoto 252 Combo score computed against the query amongst all conformations of the sampled molecule. We 253 define a hit as a sampled molecule  $m_{\text{gen}}$  with  $\mathsf{TC}(m_{\text{gen}}, m_{\text{query}}) \ge 1.2$  with respect to the query ligand 254  $m_{\text{query}}$ , in line with the heuristics used by Grebner et al. (2020). We evaluate on four metrics: number 255 of hits per query molecule, number of hits per query molecule containing a unique scaffold, max Tanimoto Score per query molecule and number of query molecules with at least one hit. 256

257 We then evaluate our model as well as the Pharmacophore Guided Molecular Generation (PGMG) 258 model proposed by Zhu et al. (2023) for pharmacophore-based generation. The PGMG model is based 259 on a conditional variational autoencoder (Kingma and Welling, 2022; Sohn et al., 2015) architecture. 260 At training time, the PGMG model learns to denoise a corrupted SMILES string conditioned on a 261 pharmacophore point cloud that is computed from the uncorrupted SMILES string. At inference time, the model generates SMILES strings corresponding to a given pharmacophore point cloud. The point 262 cloud is encoded using a Gated Graph Convolutional Neural Network (Bresson and Laurent, 2018) that 263 is equivariant to rotation, and a transformer (Vaswani et al., 2017) architecture is used to process the 264 encoded representation along with the SMILES string. For the ChEMBL dataset, we use the pretrained 265 model provided by Zhu et al. (2023). For GEOM-drugs, we train a model using their codebase and 266 use the best checkpoint output by their training script. 267

For each of the candidate molecules, we compute the pharmacophore-shape profile  $\xi$  and pass it as input to VoxCap. We sample several SMILES strings from  $p(\mathbf{s}|\xi;\theta,\tau)$  using different values of  $\tau$ as well as top-k sampling (Fan et al., 2018). For the PGMG model, we pass the SMILES string of

270	Method	Dataset	Hits	Unique Scaffold Hits	Max	# Queries with $\geq 1$ hit
271	Dataset baseline		0	0	1.16	43
272	PGMG	GEOM-drugs	11.5	7	1.44	77
273	VoxCap	-	116.5	55.5	1.77	97
274	Dataset baseline		0	0	1.12	32
275	PGMG	ChEMBL	11.5	8.5	1.39	77
276	VoxCap		115	72	1.77	95

Table 1: De-novo generation where the queries are drawn from a test set with the same distribution as the training set. Values for hits, unique scaffold hits and max score are medians across the 100 query molecules.

Method	Training/Test Dataset	Hits	Unique Scaffold Hits	Max	# Queries with $\geq 1$ hit
PGMG	ChEMPL/CEOM days	12	8	1.40	73
VoxCap	CHEMBL/GEOM-drugs	85.5	48.5	1.65	90
PGMG	GEOM drugs/ChEMRI	1	1	1.23	57
VoxCap	OLOM-drugs/CILLWIDL	23.5	12	1.47	77

Table 2: De-novo generation where the queries are drawn from a test set with a different distribution than the training set. Values for hits, unique scaffold hits and max score are medians across the 100 query molecules.

289 290

287

288

277

278

279

280 281

283 284 285

291

292 the query molecule into the model along with a pharmacophore graph computed from this string using the code provided by the authors, and sample several SMILES strings. Before evaluating, we 293 remove invalid and duplicate generated molecules. We also remove generated molecules that are identical to the query molecule, as the purpose of this task is to generate novel, diverse molecules with 295 a similar pharmacophore-shape profile to the reference molecule, not to reproduce the query molecule. 296 For each model, we select the first 1000 valid, non-duplicate generated molecules corresponding 297 to each candidate molecule, then obtain conformations and compute metrics for these molecules in 298 the same way as for the dataset baseline. (For VoxCap, before selection, the generated molecules are 299 arranged in ascending order of the  $\tau$  value used to generate them, so that we favour selecting molecules 300 sampled from the distribution on which the model was trained, where  $\tau = 1$ .) In this way, we compare 301 approaches using the same budget for ROCS evaluation, which is the time-consuming part of the 302 evaluation, compared to forward passes of neural networks which are relatively cheap.

We use the validation set for early stopping and model selection based on the total number of unique scaffold hits in a small set of 100 generated samples for each of 10 randomly selected validation set query molecules.

Results on GEOM-drugs and ChEMBL are given in Table 1. (See also Supplementary Table 1 for mean 307 and standard deviation summary statistics.) Metrics corresponding to the dataset baseline are very 308 low, showing that two randomly selected molecules from the same marginal distribution are expected 309 to have low overlap. VoxCap outperforms PGMG by large margins across all metrics on both datasets. 310 In contrast to PGMG, VoxCap produces at least one hit for almost every query molecule in each test 311 set. Of note, along with the high number of hits, is the high number of unique scaffold hits, showing 312 that the model produces a diverse set of hypotheses with matching pharmacophore-shape profiles, 313 providing diverse avenues for ligand exploration corresponding to a given query. The top score among 314 all generated ligands, averaged over all query molecules, is also higher for VoxCap than for PGMG 315 across both datasets. VoxCap is therefore able to generate hits for queries drawn from the distribution on which it was trained. This in-distribution evaluation relates to drug-discovery campaigns where 316 lead compounds are similar to library compounds that the campaign has access to, or has produced 317 before, which have been used to train the model. 318

The mean and median values for all VoxCap metrics are comparable in all cases, showing that VoxCap's successful performance is not merely due to an extremely high contribution to the metrics from a small number of queries and low contribution from the rest, but is instead due to high contributions from many queries. The standard deviation across the 100 queries is, however, high for both the hits and unique scaffold hits metrics. Bringing this number down would correspond to more predictable performance for a given query and is an important area of future work. 324 We similarly measure out-of-distribution performance in Table 2 and Supplementary Table 2. For this 325 task, models trained on GEOM-drugs are evaluated on query molecules from the ChEMBL subset, and 326 vice-versa. We see that VoxCap once again outperforms PGMG across all metrics by large margins. 327 VoxCap is hence able to generate hits for queries drawn from a distribution that is different to the one on 328 which it was trained. This relates to drug-discovery campaigns where lead compounds are significantly different from compounds that the campaign has access to, or has produced before. Results from 329 models trained on ChEMBL are better than the corresponding models trained on GEOM-drugs, 330 possibly because of the larger number of molecules in ChEMBL. 331

332 333

367

371

#### 4.2 FAST MOLECULAR LIBRARY SEARCH

334 Carrying out a brute-force ROCS search against a given database (also known as a molecular library) 335 for each new query molecule in an experimental setting is a time-consuming and costly process. 336 We propose an inexpensive approximation to this workflow by taking advantage of the generative 337 capability and cheap inference cost of our model. Given a query ligand m from the test set, we compute 338 its pharmacophore-shape profile  $\xi$ , use VoxCap to generate several sample SMILES strings s from  $p(\mathbf{s}|\boldsymbol{\xi};\boldsymbol{\theta},\tau)$ , and then remove duplicate and invalid molecules. For each of these  $n_q$  generated molecules 340 we then find, in the union of the train and validation sets, the  $n_a$  closest 2D analogs by Tanimoto 341 similarity (i.e. Jaccard similarity) of Morgan fingerprints (a bit vector featurisation of a molecule, 342 each element of which indicates whether a particular substructure is contained in the molecular graph), and compute the ROCS score for each of these analogs against the query molecule. A diagram of this 343 algorithm is shown in Figure 2. 344



Database molecules with  $\mathsf{TC} \geq 1.2$  are hits

Figure 2: *Fast search* workflow, where we focus on the procedure applied to one of  $n_g$  generated unique, valid molecules. Fingerprints of database molecules can be precomputed and reused for each new query to further reduce computational cost.

Our proposed approach drastically reduces the cost of database search. The 2D search we employ is fast and inexpensive, with negligible computational cost compared with 3D conformation generation, alignment to the query molecule and Tanimoto Combo score computation. Our approach brings down the number of 3D comparisons from O(database size) to  $O(n_g \times n_a)$ , a difference of several orders of magnitude depending on database size. As an example, the Enamine Real and Enamine Diverse databases consist of 60 billion compounds and 100 million compounds respectively, whereas  $n_g \times n_a$  is typically on the order of 1000. Furthermore, running brute force 3D searches against the full Enamine Real database is intractable, especially in a generative design framework where we wish to run searches
for hundreds or thousands of generated molecules, whereas our fast search workflow is feasible. Our approach is advantageous to simply running the 2D search over the entire database for each query molecule
as we are interested in finding compounds that are diverse in substructure space but similar to the query
in 3D pharmacophore-shape space; not analogs that are simply similar to the query in substructure space.

We evaluate the accuracy of this workflow as follows. For each of 11 randomly selected query 384 molecules from the ChEMBL test set, we run a ROCS search against every molecule in the union 385 of the train and validation sets excluding the query molecule. We call this the brute force approach. 386 We also use VoxCap to generate several molecules given the query molecule's pharmacophore-shape 387 profile, and find either (i)  $n_a = 5$  analogs for each of up to  $n_a = 100$  unique, valid molecules, or (ii) 388  $n_a = 1$  analog for each of up to  $n_a = 500$  molecules. We compute these analogs' ROCs scores against the query molecule as described above. We then compute the number of hits in this set of analogs 389 and compare with the number of hits in the set of top 500 molecules by ROCS score from the brute 390 force approach for each query molecule. Results are given in Table 3. 391

392 We find that this workflow yields a significant number of hits and unique scaffold hits; the two settings 393 of  $n_a$  and  $n_a$  yield similar results. The number of ROCS comparisons per query (up to a constant,  $\alpha$ , the 394 average number of conformers per database molecule) is  $n_a \times n_a$ . This is several orders of magnitude 395 below that of the brute force search, which is the size of the database. This difference would hence be even more pronounced in a real-world drug discovery application for databases such as Enamine 396 Real/Diverse. We would also expect to find a greater number of 2D matches and hits from VoxCap-397 generated molecules for these databases due to their larger size. Our workflow hence enables the search 398 of large-scale databases for sets of diverse hits, a current bottleneck in the drug design pipeline. 399

400 Our results show that there are limitations to our approach, namely the number of hits is over an 401 order of magnitude lower than the brute force method, suggesting significant room for improvement. We explore why this may be the case. We observe that the highest 2D Tanimoto similarity for 402 any given generated molecule is generally low, with a median top-1 similarity of 0.38 for the 403  $n_a = 500, n_a = 1$  setting. Hence many of the generated molecules do not have close 2D analogs in 404 the dataset. Supplementary Figure 2 shows a moderate correlation (Pearson r = 0.4) between 2D and 405 3D similarities. The lack of close 2D analogs for most of the generated molecules could therefore 406 be a major reason for the low number of hits relative to the brute force approach. We also observe that 407 there is a significant number of duplicate 2D hits per query, where the 2D search procedure maps on 408 average 2.4 generated molecules to the same molecule in the database for the  $n_a = 500, n_a = 1$  setting, 409 thereby decreasing the number of hits that can be found on running ROCS.

410 We also observe from our experiments that our early stopping criterion selects a checkpoint after only 411 a few epochs of training. Further training causes performance on sampling metrics to drop even as 412 the validation cross-entropy loss continues to decrease. We hypothesise that the input representation 413 is too specific to each individual molecule, so with continued training the model learns a distribution 414 that is sharply peaked around the most likely molecule without adequately learning other modes of 415 the distribution. This could mean that, at the early stopping checkpoint, when the number of training 416 epochs completed is not large, the model has not been trained long enough to adequately capture the 417 2D structure of molecules in the dataset.

Future directions of work could involve regularisation/early stopping strategies that yield molecules
 more structurally similar to those in the dataset, as well as input representations that are smooth or
 noisy enough to facilitate the learning of several modes of the distribution over a larger number of
 training steps.

422 423 424

425 426 427

	Method	$n_g$	$\mid n_a$	Hits	Unique	Max	Queries w/ $\geq$ 1 hit	ROCS comparisons/query
-	Brute force	-	-	450	151	1.67	11	$\alpha \times 1.3 \times 10^{6}$
	VoxCap	100	5	10	9	1.58	10	$\alpha \times 500$
	VoxCap	500	1	10	8	1.58	11	$\alpha \times 500$

Table 3: Fast database search results. Unique stands for unique scaffold hits. ROCS comparisons/query are given including the constant  $\alpha$ , which is the average number of conformers per database molecule. Values for hits, unique scaffold hits and max score are medians across the 11 query molecules.

431

## 432 5 CONCLUSION

433 434

In this work, we propose the generation of molecules conditioned on voxelised 3D pharmacophore-435 shape profiles. We introduce VoxCap, a generative model that generates diverse SMILES strings 436 from these pharmacophore-shape profiles. We propose a benchmark task for *de-novo* generation 437 conditioned on pharmacophore-shape profiles and find that VoxCap outperforms the existing state of the art by large margins across all metrics. We also implement a generative model-based fast search 438 workflow to practically enable the search of large libraries for structurally diverse 3D hits against a 439 given query. We find that VoxCap, paired with this workflow, returns hits from the database for each 440 query molecule, but that there is significant room for improvement. We identify areas of future work 441 corresponding to coarser pharmacophore-shape representations and regularisation, that would enable 442 the efficient search of large-scale databases for hit compounds in drug discovery campaigns. 443

444 445

454

#### References

- Chayan Acharya, Andrew Coop, James Polli, and Alexander MacKerell. Recent advances in ligand-based drug design: Relevance and utility of the conformationally sampled pharmacophore approach. *Current computer-aided drug design*, 2011.
- 450 Simon Axelrod and Rafael Gomez-Bombarelli. Geom, energy-annotated molecular conformations
   451 for property prediction and molecular generation. *Scientific Data*, 2022.
- <sup>452</sup> Denzil Bernard, Andrew Coop, and Alexander D MacKerell. Conformationally sampled <sup>453</sup> pharmacophore for peptidic  $\delta$  opioid ligands. *Journal of medicinal chemistry*, 2005.
- 455 Xavier Bresson and Thomas Laurent. Residual gated graph convnets, 2018.
- 456 Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking
   457 models for de novo molecular design. *Journal of chemical information and modeling*, 2019.
- Cheng Chang and Peter W Swaan. Computational approaches to modeling drug transporters. *European Journal of Pharmaceutical Sciences*, 2006.
- <sup>461</sup> Angela Fan, Mike Lewis, and Yann N. Dauphin. Hierarchical neural story generation. *CoRR*, 2018.
- Pablo M González, Chayan Acharya, Alexander D MacKerell, and James E Polli. Inhibition requirements of the human apical sodium-dependent bile acid transporter (hasbt) using aminopiperidine conjugates of glutamyl-bile acids. *Pharmaceutical research*, 2009.
- Christoph Grebner, Erik Malmerberg, Andrew Shewmaker, Jose Batista, Anthony Nicholls, and Jens
  Sadowski. Virtual screening in the cloud: How big is big enough? *Journal of Chemical Information and Modeling*, 2020.
- Thomas A. Halgren. Merck molecular force field. I. basis, form, scope, parameterization, and performance of MMFF94. *J. Computational Chemistry*, 1996.
- Paul C. D. Hawkins, A. Geoffrey Skillman, and Anthony Nicholls. Comparison of shape-matching
   and docking as virtual screening tools. *Journal of Medicinal Chemistry*, 2007.
- 474 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 1997.
- 476 Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022.
- 477
   478
   479
   Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models, 2014.
- Yasuhisa Kurogi and Osman F Guner. Pharmacophore modeling and three-dimensional database
   searching for drug design using catalyst. *Current medicinal chemistry*, 2001.
- 482
   483
   484
   484
   484
   484
   485
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
   484
- 485 Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.

486 487 488	Guang Li, Linchao Zhu, Ping Liu, and Yi Yang. Entangled transformer for image captioning. In <i>ICCV</i> , 2019.
489 490	Gilda H Loew, Hugo O Villar, and Ibon Alkorta. Strategies for indirect computer-aided drug design. <i>Pharmaceutical research</i> , 1993.
491 492 493	Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, and Alan L. Yuille. Explain images with multimodal recurrent neural networks, 2014.
494 495	JS Mason, AC Good, and EJ Martin. 3d pharmacophores in drug discovery. <i>Current pharmaceutical design</i> , 2001.
496 497 498	Alan D McNaught, Andrew Wilkinson, et al. <i>Compendium of chemical terminology</i> . Blackwell Science Oxford, 1997.
499 500 501	David Mendez, Anna Gaulton, A Patrícia Bento, Jon Chambers, Marleen De Veij, Eloy Félix, María Paula Magariños, Juan F Mosquera, Prudence Mutowo, Michał Nowotka, et al. Chembl: towards direct deposition of bioassay data. <i>Nucleic acids research</i> , 2019.
502 503 504	OpenEye, Cadence Molecular Sciences. Rocs 3.6.2.0, 2024. URL http://www.eyesopen.com. Santa Fe, NM.
505 506 507	Gabriele Orlando, Daniele Raimondi, Ramon Duran-Romaña, Yves Moreau, Joost Schymkowitz, and Frederic Rousseau. Pyuul provides an interface between biological structures and deep learning algorithms. <i>Nature communications</i> , 2022.
508 509 510	RS Pearlman. 3d qsar in drug design. theory, methods, and applications. 3D QSAR in drug design-theory, methods, and applications, 1993.
511 512 513	Pedro O Pinheiro, Joshua Rackers, Joseph Kleinhenz, Michael Maser, Omar Mahmood, Andrew Watkins, Stephen Ra, Vishnu Sresht, and Saeed Saremi. 3d molecule generation by denoising voxel grids, 2023.
514 515 516	Pedro O Pinheiro, Arian Jamasb, Omar Mahmood, Vishnu Sresht, and Saeed Saremi. Structure-based drug design by denoising voxel grids, 2024.
517 518 519	Kristina Preuer, Philipp Renz, Thomas Unterthiner, Sepp Hochreiter, and Gunter Klambauer. Fréchet chemnet distance: a metric for generative models for molecules in drug discovery. <i>Journal of chemical information and modeling</i> , 2018.
520 521 522	Matthew Ragoza, Tomohide Masuda, and David Ryan Koes. Learning a continuous representation of 3D molecular structures with deep generative models. In <i>Neurips, Structural Biology workshop</i> , 2020.
523 524 525	Matthew Ragoza, Tomohide Masuda, and David Ryan Koes. Generating 3d molecules conditional on receptor binding sites with deep generative models. <i>Chemical science</i> , 2022.
526 527	Bharath Ramsundar, Peter Eastman, Patrick Walters, Vijay Pande, Karl Leswing, and Zhenqin Wu. <i>Deep Learning for the Life Sciences</i> . O'Reilly Media, 2019.
529 530	Sereina Riniker and Gregory A. Landrum. Better informed distance geometry: Using what we know to improve conformation generation. <i>Journal of Chemical Information and Modeling</i> , 2015.
531 532	Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.
533 534 535 536	Thomas S. Rush, J. Andrew Grant, Lidia Mosyak, and Anthony Nicholls. A shape-based 3-d scaffold hopping method and its application to a bacterial proteinprotein interaction. <i>Journal of Medicinal Chemistry</i> , 2005.
537 538	Saeed Saremi and Aapo Hyvärinen. Neural empirical Bayes. JMLR, 2019.
	$M^{(1)}_{1} = 0 + 1^{(1)}_{1} + 1^{(1)}_{2} = 0 + 1^{(1)}_{1} + 0^{(1)}_{1} + 1^{(1)}_{2} = 0 + 1^{(1)}_{2} + 0^{(1)}_{2} + 0^{(1)}_{2} = 0$

539 Miha Skalic, José Jiménez, Davide Sabbadin, and Gianni De Fabritiis. Shape-based generative modeling for de novo drug design. *Journal of chemical information and modeling*, 2019.

540 541 542	Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In <i>Advances in Neural Information Processing Systems</i> , 2015.
543 544	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017.
545 546	Clement Vignac, Nagham Osman, Laura Toni, and Pascal Frossard. Midi: Mixed graph and 3d denoising diffusion for molecule generation. In <i>ECML</i> , 2023.
547 548 549	Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator., 2015.
550 551	David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. <i>Journal of Chemical Information and Computer Sciences</i> , 1988.
552 553 554	Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention, 2016.
555 556 557 558 559	Huimin Zhu, Renyi Zhou, Dongsheng Cao, Jing Tang, and Min Li. A pharmacophore-guided deep learning approach for bioactive molecular generation. <i>Nature Communications</i> , 2023.
560 561	
562 563	
564 565	
566	
568	
569 570	
571 572	
573 574	
575 576	
577	
578	
580 581	
582 583	
584 585	
586	
588	
589 590	
591 592 593	

## <sup>594</sup> A MODEL ARCHITECTURE



# 648 B VOXELISATION

650 We represent pharmacophores and molecular shape as voxelised densities. Given an input molecular 651 conformation, we compute the locations of all pharmacophores in the molecule, resulting in a 652 point cloud for each of the 6 pharmacophore types. We also compute a shape point cloud, whose 653 points' locations are the same as those of the molecule's atoms. We now have a set of 7 point clouds 654  $C = \{C_c\}_{c=1}^7$ . The *c*th point cloud is given by  $C_c = \{a_{c,n}\}_{n=1}^{N_c}$  where  $a_{c,n}$  is the *n*th point in point cloud 655  $C_c$  and  $N_c$  is the number of points in point cloud  $C_c$ .

We voxelise the pharmacophore and shape point clouds in a similar manner to how Pinheiro et al. (2023) voxelise atomic densities. First we convert each point  $a_{c,n}$  into a 3D Gaussian-like density:

$$V_{a_{c,n}}(d,r) = \exp\left(-\frac{d^2}{(.93 \cdot r)^2}\right),$$
(4)

where  $V_{a_{c,n}}$  is defined as the fraction of occupied volume by point  $a_{c,n}$  at distance d from its center. We use radius r = 1 for all points. We then compute the occupancy of each voxel in the grid at each of the grid points (i,j,k) at each channel c:

$$\operatorname{Occ}_{c,i,j,k} = 1 - \prod_{n=1}^{N_c} \left( 1 - V_{a_{c,n}} ( \| C_{i,j,k} - x_{a_{c,n}} \|, r ) \right),$$
(5)

where  $C_{i,j,k}$  are the coordinates (i,j,k) in the grid and  $x_{a_{c,n}}$  is the coordinates of point  $a_{c,n}$ . Since we use c to index both channels and point clouds, points in point cloud  $C_c$  are assigned to channel c. Each voxel in the grid has a value between 0 and 1. We use the python package PyUUL (Orlando et al., 2022) to generate the voxel grids.

#### C DATA AUGMENTATION

During training, we randomly augment each training sample. This augmentation comprises a random translation by a value uniformly sampled from [-1,1] for each axis, and a random rotation by a value uniformly sampled from  $[0, 2\pi)$  for each of the three Euler angles. This augmentation is applied to the point clouds before voxelisation.

#### D DE-NOVO GENERATION

	Method	Dataset	Hits	Unique Scaffold Hits	Max
-	Dataset baseline	GEOM-drugs	$1.37\pm2.88$	$1.33\pm2.75$	$1.18\pm0.17$
	PGMG		$20.14\pm23.76$	$10.01 \pm 10.96$	$1.41 \pm 0.30$
	VoxCap		$126.08\pm98.22$	$73.12\pm63.44$	$1.69\pm0.24$
-	Dataset baseline	ChEMBL	$0.63 \pm 1.25$	$0.62 \pm 1.25$	$1.12\pm0.15$
	PGMG		$47.98 \pm 74.71$	$27.56\pm37.83$	$1.42\pm0.27$
	VoxCap		$150.70 \pm 121.50$	$90.69 \pm 80.00$	$1.70\pm0.25$

Supplementary Table 1: De-novo generation where the queries are drawn from a test set with the same distribution as the training set. Values for hits, unique scaffold hits and max score are means  $\pm$  standard deviations across the 100 query molecules.

702	Method	Training/Test Dataset	Hits	Unique Scaffold Hits	Max
703	PGMG	ChEMPL/CEOM drugs	$51.21 \pm 90.04$	$27.73 \pm 43.86$	$1.39\pm0.30$
704	VoxCap	CHEMBL/GEOWI-drugs	$111.29 \pm 100.01$	$67.56\pm65.54$	$1.61\pm0.28$
705	PGMG	CEOM drage/ChEMPI	$11.97\pm20.16$	$5.97 \pm 9.93$	$1.26\pm0.24$
706	VoxCap	GEOM-ulugs/CIIEMBL	$55.39\pm75.62$	$33.63 \pm 49.02$	$1.46\pm0.30$
707					

Supplementary Table 2: De-novo generation where the queries are drawn from a test set with a different distribution than the training set. Values for hits, unique scaffold hits and max score are means  $\pm$  standard deviations across the 100 query molecules.

#### E FAST DATABASE SEARCH

Method $n_g$		$ n_a $	Hits	Unique Scaffold Hits	Max
Brute force	-	-	$326.27 \pm 214.53$	$158.82 \pm 121.47$	$1.64\pm0.21$
VoxCap	100	5	$20.64\pm17.55$	$12.73\pm9.31$	$1.52\pm0.55$
VoxCap	500	1	$16.09\pm10.63$	$11.82\pm7.68$	$1.65\pm0.21$

Supplementary Table 3: Fast database search results. Values for hits, unique scaffold hits and max score are means  $\pm$  standard deviations across the 11 query molecules



Supplementary Figure 2: 3D vs 2D similarities, from fast database search experiment with  $n_g = 500, n_a = 1$