

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 UNIFIED CROSS-SCALE 3D GENERATION AND UN- DERSTANDING VIA AUTOREGRESSIVE MODELING

Anonymous authors

Paper under double-blind review

## ABSTRACT

3D structure modeling is essential across scales, enabling applications from fluid simulation and 3D reconstruction to protein folding and molecular docking. Yet, despite shared 3D spatial patterns, current approaches remain fragmented, with models narrowly specialized for specific domains and unable to generalize across tasks or scales. We propose Uni-3DAR, a unified autoregressive framework for cross-scale 3D generation and understanding. At its core is a coarse-to-fine tokenizer based on octree data structures, which compresses diverse 3D structures into compact 1D token sequences. We further propose a two-level subtree compression strategy, which reduces the octree token sequence by up to 8x. To address the challenge of dynamically varying token positions introduced by compression, we introduce a masked next-token prediction strategy that ensures accurate positional modeling, significantly boosting model performance. Extensive experiments across multiple 3D generation and understanding tasks, including small molecules, proteins, polymers, crystals, and macroscopic 3D objects, validate its effectiveness and versatility. Notably, Uni-3DAR surpasses previous state-of-the-art diffusion models by a substantial margin, achieving up to 256% relative improvement while delivering inference speeds up to 21.8x faster.

## 1 INTRODUCTION

3D structure modeling underpins a wide range of real-world applications, spanning the planetary-scale dynamics of celestial bodies to the angstrom-scale arrangements of atoms and electrons. At the macroscopic level, it enables 3D object reconstruction, computational fluid dynamics simulations, and climate forecasting; at the microscopic level, it supports protein structure prediction (Jumper et al., 2021), crystal generation (Jiao et al., 2023), molecular dynamics (Wang et al., 2018a), and molecular docking (Alcaide et al., 2024).

Despite these shared spatial principles, 3D modeling tasks have largely evolved in silos. Models tailored for macroscopic structures fail to transfer to microscopic domains, and even applications at similar scales rarely generalize. For instance, a model designed for crystal generation cannot be directly applied to protein folding (Xie et al., 2021; Jiao et al., 2023). This fragmented development hinders data reuse and results in redundant, highly specialized models rather than a unified solution.

To overcome this fragmentation, we propose Uni-3DAR, a unified autoregressive framework for cross-scale 3D generation and understanding. At its core is a tokenizer that efficiently compresses diverse 3D structures into discrete 1D token sequences. Leveraging these compressed sequences, our autoregressive model unifies generative and understanding tasks within a single architecture.

The proposed tokenizer uses an octree data structure to compress the full-size 3D grid both losslessly and efficiently. As illustrated in Fig. 2 (a) and (b), we construct an octree by recursively subdividing the space up to a maximum depth of  $L$ . To adapt to data sparsity, branches corresponding to empty regions are pruned, resulting in a maximum of  $8^{L-1}$  leaf grid cells (but most will be pruned due to sparsity). We then introduce a fine-grained tokenization that encodes details within each occupied leaf cell (we call it a “3D patch”), such as atomic types and precise coordinates for molecules, or more general VQVAE tokens (Van Den Oord et al., 2017). Concatenating these tokens level by level produces a hierarchical, coarse-to-fine 1D token sequence that effectively represents the 3D structure (fig. 2(c)). Furthermore, we compress each two-level subtree (eight subcells) into a single 8-bit token instead of assigning an individual occupancy token to each node (Fig. 2 (d)). Since each

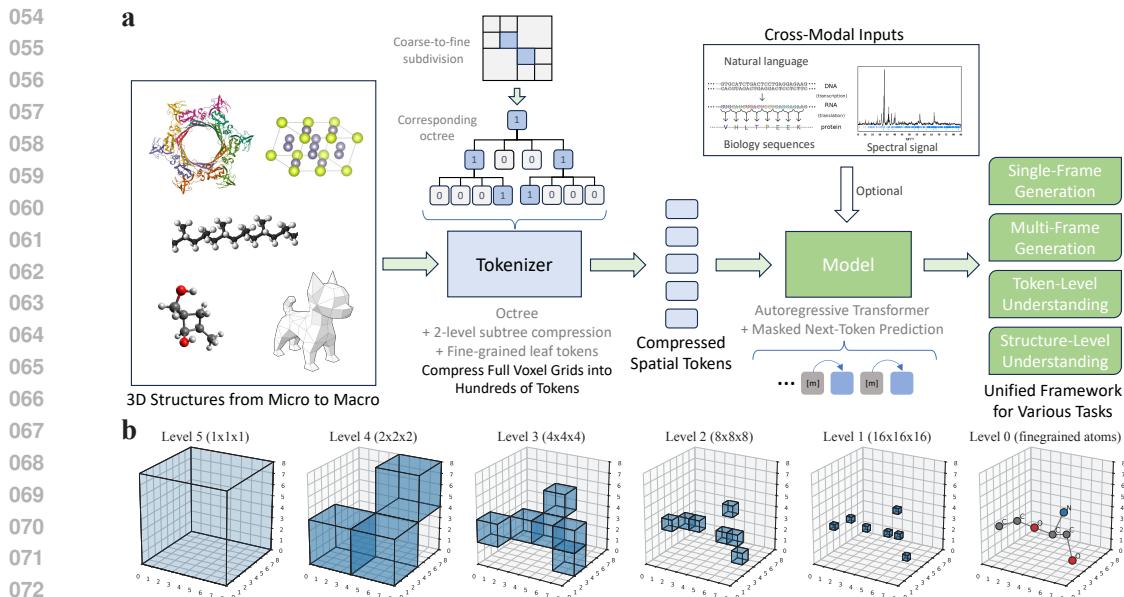


Figure 1: **Uni-3DAR Overview.** (a) A coarse-to-fine octree-based tokenizer converts 3D structures into 1D sequences (details in fig. 2). The tokens are modeled by an autoregressive transformer trained with masked next-token prediction (details in fig. 3) and can be optionally conditioned on cross-modal inputs (e.g., text, biological sequences, spectra). A single model supports single- and multi-frame generation as well as token- and structure-level understanding. (b) An example of octree from coarse level to fine level. Uni-3DAR generates tokens in a coarse-to-fine order: high-level occupancy tokens followed by level-0 tokens that capture local details (e.g., atom types and coordinates). The merits of octree over other 3D representations are discussed in Appendix A.

subcell is binary (empty or not), grouping eight subcells yields  $2^8 = 256$  distinct states, reducing the sequence length approximately 8x and converting 8 binary classifications into one 256-class task.

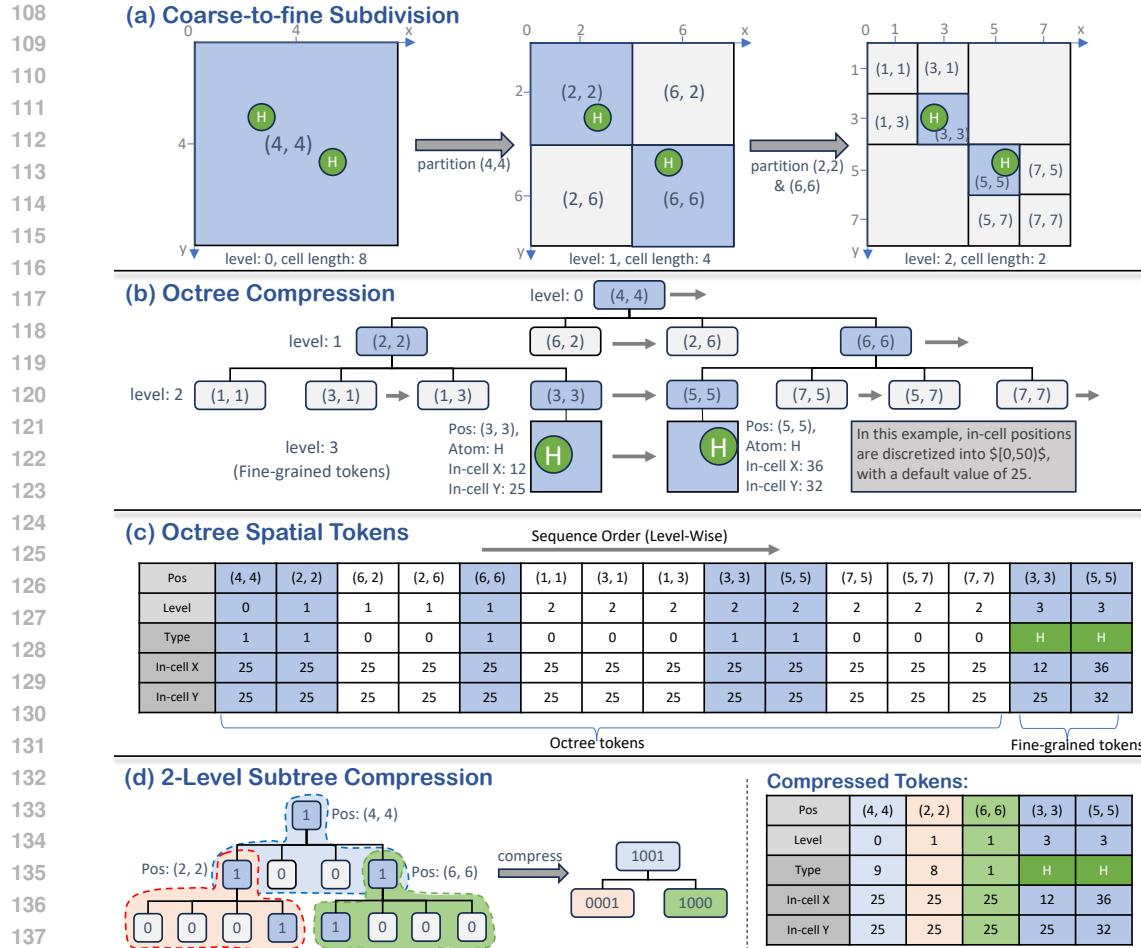
However, the octree compression with empty tokens pruned disrupts the spatial mapping, meaning adjacent tokens no longer correspond to uniform intervals in the original 3D space. Unlike in 2D images with fixed patch positions, the model cannot reliably predict the next token without knowing its explicit target coordinates. We found that simply appending the next position to the current token yielded unsatisfactory results. To address this challenge, we propose a *masked next-token prediction* strategy. As illustrated in Fig. 3 (a), our method duplicates each token so that it appears twice with the same positional embedding. We then replace the first copy with a [MASK] token. The model still performs next-token prediction, but the prediction is made exclusively at the masked position. This setup ensures that the prediction is conditioned on the correct positional information of the intended target token, effectively resolving the issue of dynamic token positions. Although this approach doubles the sequence length, it achieves significant performance gains as validated in appendix D.2.

Uni-3DAR is built on several technical innovations: (1) a **coarse-to-fine octree-based tokenization** for efficient representation, (2) a **2-level subtree compression** to reduce sequence length, (3) a **unified fine-grained structural representation** (for “3D patch”) to capture local details, and (4) a **masked next-token prediction** strategy to handle dynamic token positions, which enable our key contributions:

1. *Unified Cross-Scale 3D Modeling.* Leveraging the proposed coarse-to-fine tokenizer, Uni-3DAR can process a wide range of 3D structures, from macroscopic to microscopic scales.

2. *Unified Generation and Understanding.* Uni-3DAR seamlessly unifies 3D structural generation and understanding tasks within a single framework. As illustrated in Fig. 3 (b), different tasks use distinct tokens, ensuring clear separation without interference.

3. *High Efficiency.* Thanks to the octree and two-level subtree compression, Uni-3DAR represents 3D structures with far fewer tokens. For example, while O Pinheiro et al. (2023) requires  $32^3 = 262,144$  tokens for a small molecule, Uni-3DAR needs only hundreds, and can scale to large proteins with thousands of atoms with deeper octree levels (section 3.4). Moreover, appendix D.3 shows that Uni-3DAR is approximately 21.8x faster than prior diffusion-based models.



138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149

Figure 2: Overview of Uni-3DAR tokenization (illustrated in 2D using quadtree for clarity). **(a)** Adaptive coarse-to-fine subdivision of grid cells, where darker nodes indicate non-empty cells that can be further partitioned. **(b)** This partitioning process constructs an octree, providing a lossless compression of the full-size 3D grid. **(c)** Uni-3DAR’s tokenization consists of two components: hierarchical spatial compression via an octree and fine-grained structural tokenization. Each node’s position is determined by its tree level and cell center. **(d)** The proposed 2-level subtree compression reduces the octree tokens by 8x (4x in the illustrated quadtree).

150  
151  
152  
153  
154  
155  
156  
157  
158  
159  
160  
161

**4. High Accuracy.** Extensive experiments across diverse tasks—including macroscopic 3D shape generation (table 4), molecular (table 1), crystal generation (table 2), protein pocket prediction (table 6), molecular docking (table 7), and molecular pretraining (tables 8 and 9)—demonstrate Uni-3DAR’s superior or competitive performance compared to existing methods. Notably, Uni-3DAR consistently outperforms diffusion-based models. Ablation studies (table 12) highlight the benefits of unifying generation and understanding and validate the effectiveness of each component.

## 2 METHOD

### 2.1 DYNAMIC COARSE-TO-FINE TOKENIZATION FOR 3D STRUCTURES

153  
154  
155  
156  
157  
158  
159  
160  
161

3D structures are inherently sparse: at the microscopic scale, most space is empty except for scattered atoms; at the macroscopic level, detailed representations are only needed at object surfaces, with most volume remaining empty. Using a full-size voxel grid is thus highly inefficient. To address this, we propose a hierarchical, coarse-to-fine tokenization of 3D structures that exploits this sparsity. As shown in Fig. 1, our approach consists of two parts: (1) a hierarchical compression of 3D space using an octree, and (2) a fine-grained structural tokenization.

The first component is the octree, an efficient data structure for lossless 3D grid compression. Starting with a single cell covering the entire structure, we recursively subdivide it: if a cell contains atoms, it is partitioned further. Each subdivision halves each dimension, producing  $2^3 = 8$  equal subcells (hence “octree”). This process continues for  $L$  levels. If  $c_0$  is the root cell length, the

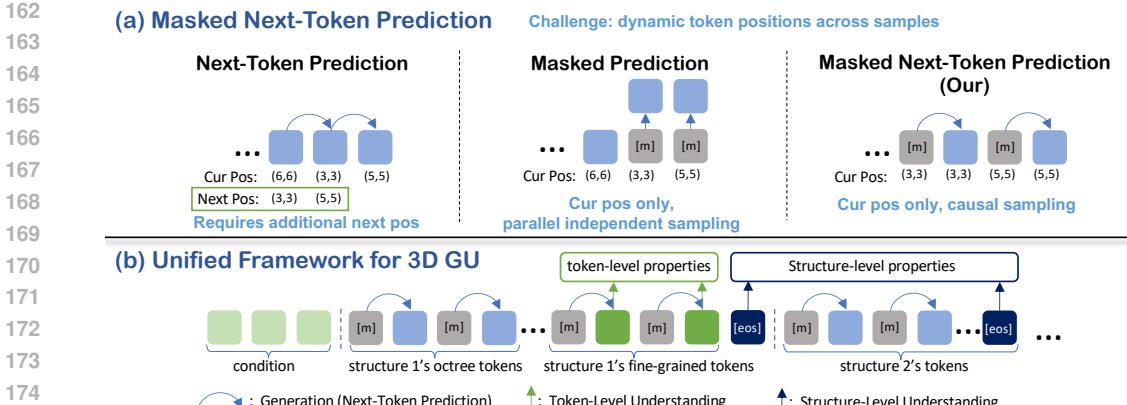


Figure 3: (a) **Masked Next-Token Prediction.** To handle the challenge of dynamically positioned tokens in sparse 3D structures, Uni-3DAR decouples position and content generation. Unlike standard next-token prediction, we first infer the next token’s position from the octree hierarchy, place a “[MASK]” token, and then have the model predict only its content (e.g., occupancy or fine-grained properties). (b) **Unified Framework for 3D Generation and Understanding.** The Uni-3DAR architecture is a versatile, multi-task model. It supports autoregressive generation of complex 3D structures (blue arrows) and can be prompted to perform both token-level (green arrows) and structure-level (blue box) understanding tasks within a single framework.

cell length at level  $L - 1$  is  $c_{L-1} = c_0/2^{L-1}$ . We refer to these leaf subcells as fine-grained “3D patches,” which are then tokenized as detailed in the following paragraph.

The second component is the fine-grained tokenization of structural details. While the octree effectively identifies coarse, non-empty regions, it lacks finer details such as atom types and precise coordinates (microscopic) or surface features (macroscopic). Although using deeper octrees can capture more detail (Ibing et al., 2023), this approach becomes inefficient due to the rapidly increasing number of tokens. Instead, inspired by the use of 2D image patches (Alexey, 2020), we treat the contents of each final-level non-empty region as a “3D patch.” These patches can be processed in various ways; for instance, they can be quantized into discrete tokens for autoregressive prediction, similar to VQ-VAE (Van Den Oord et al., 2017), or modeled using a patch-level diffusion loss for continuous vector representations (Li et al., 2024b) (ablation in Table 12). In our experiments, we demonstrate this flexibility by using raw atom types and coordinates as fine-grained tokens for microscopic data (we set the patch size to ensure each 3D patch only contains one atom), and VQ-VAE tokens for macroscopic data. More details are in Appendix B.

Finally, we concatenate tokens level by level. Beyond token content, we represent each token’s positional information using its tree level and the spatial coordinates of its cell center. For instance, the root cell is at level 0 with a center at  $(c_0/2, c_0/2, c_0/2)$ . During autoregressive prediction, since octree tokens are dynamically unfolded level by level, the positions of all tokens at the current level are known based on the predictions from previous level. This explicit knowledge of the token position is crucial, as autoregressive models predict only token content.

**2-Level Subtree Compression** Although octree tokenization avoids cubic cell growth, it remains inefficient for large 3D structures. Each level has up to  $8N$  tokens ( $N$  = non-empty final-level cells), totaling up to  $8NL$  tokens across  $L$  levels, about two orders of magnitude larger than  $N$ . To reduce this, we introduce 2-level subtree compression, merging a parent and its 8 children into a single token. As the parent’s type is always 1, the subtree is fully represented by its 8 children’s types, yielding  $2^8 = 256$  possible states. This cuts token count by  $8\times$ , down to at most  $N(L - 1)$  tokens. For positional information of the compressed nodes, we retain the their parent’s center and level.

## 2.2 MASKED NEXT-TOKEN PREDICTION FOR DYNAMIC TOKEN POSITIONS

In standard autoregressive models, such as those for text, token positions follow a fixed, sequential order (e.g., token  $i + 1$  always follows token  $i$ ). This static structure makes the next token’s position implicitly known, obviating the need for its explicit prediction. In contrast, our coarse-to-fine 3D tokenization generates a token sequence where positions are *dynamic* and vary across different structures. This variability introduces a significant challenge: inferring the next token’s position becomes non-trivial, making it preferable to provide this information to the model explicitly.

216 A straightforward approach is to encode both the current and next positions within each token (Ibing  
 217 et al., 2023). However, we found this method leads to suboptimal performance (Table 12). We  
 218 hypothesize that the unpredictable nature of the next token’s position introduces noise that degrades  
 219 the current token’s representation. This intuition is supported by prior work like Yan et al. (2022),  
 220 which decouples position and content prediction into two separate transformer modules.

221 Another promising direction is masked prediction (Chang et al., 2022; Li et al., 2024b), where a  
 222 model predicts the content of a masked token given its position. This has proven effective for gen-  
 223 erative tasks with non-sequential or randomized token orders (Li et al., 2024b; Pang et al., 2024).  
 224 However, directly applying conventional masked prediction to our framework is problematic. First,  
 225 it typically relies on bi-directional attention, whereas our hierarchical tokenization unfolds uni-  
 226 directionally. Second, it often requires parallel, non-causal sampling, which necessitates complex,  
 227 rule-based inference strategies to balance performance and efficiency (Li et al., 2024b).

228 To resolve these issues, we introduce **Masked Next-Token Prediction** (MNTP), a simple yet effec-  
 229 tive method that integrates masked prediction into a standard autoregressive framework. The core  
 230 idea is to duplicate each token. For a given token, we first generate a placeholder with its content  
 231 replaced by a special [MASK] symbol while retaining its correct position. This is immediately fol-  
 232 lowed by a second token at the *same position* but with the true content. The model’s objective is to  
 233 predict the content of this second token, conditioned on the [MASK] token and all preceding tokens.

234 This formulation effectively reframes next-token prediction as a masked prediction task: the model  
 235 is given a position with a mask and asked to fill in the content. This approach elegantly sidesteps  
 236 the challenge of predicting dynamic next positions. Compared to conventional masked prediction,  
 237 MNTP preserves a strictly causal, uni-directional attention flow, eliminating the need for complex  
 238 sampling schemes. While this duplication doubles the sequence length, we demonstrate in Appen-  
 239 dix D.2 that the substantial performance gains justify this trade-off. Furthermore, through tar-  
 240 geted optimizations discussed in Appendix D.3, the impact on inference latency is modest, with only  
 241 a 15%–30% increase compared to standard next-token prediction (appendix D.3).

### 242 2.3 UNIFIED 3D GENERATION AND UNDERSTANDING FRAMEWORK

244 By integrating techniques in sections 2.1 and 2.2, Uni-3DAR provides a unified framework for a  
 245 wide range of 3D tasks (Figure 3(b)). The model architecture assigns distinct roles to different  
 246 token types, enabling it to handle four primary task categories individually or jointly:

- 247 1. **Single-Frame Generation** (sections 3.1, 3.3, 3.4 and 3.6): Generating a single 3D structure,  
 248 either unconditionally or conditioned on external modalities like text or chemical properties. This  
 249 is accomplished using the masked tokens for autoregressive generation.
- 250 2. **Multi-Frame Generation** (sections 3.2 and 3.5) Autoregressively producing a sequence of 3D  
 251 structures (multiple octrees), such as a molecular dynamics trajectory, molecular docking, or  
 252 pocket-based generation. Each frame is distinguished by a unique frame-index embedding.
- 253 3. **Token-Level Understanding** (section 3.4) Predicting properties of local components (e.g.,  
 254 atomic forces or partial charges) by attaching a prediction head to the fine-grained tokens.
- 255 4. **Structure-Level Understanding** (section 3.6): Predicting global properties of an entire structure  
 256 (e.g., solubility, toxicity) via a prediction head on the final “[EoS]” token. This allows Uni-3DAR  
 257 to be pre-trained on large-scale unlabeled 3D data and efficiently fine-tuned for downstream tasks.

258 This versatile design allows for seamless joint training across these diverse tasks. Each token type  
 259 serves a clear purpose: masked tokens drive generation, fine-grained tokens facilitate local under-  
 260 standing, and the “[EoS]” token enables global understanding.

262 Furthermore, the autoregressive nature of Uni-3DAR inherently supports multi-modal conditioning,  
 263 which is critical for many scientific applications. For instance, a protein’s amino acid sequence  
 264 can guide the generation of its 3D fold. Similarly, experimental data like Powder X-ray Diffraction  
 265 (PXRD) spectra can constrain the prediction of a crystal structure, a task we explore in section 3.2.

## 266 3 EXPERIMENTS

268 We conducted extensive experiments to validate Uni-3DAR across diverse benchmarks. This section  
 269 summarizes the key findings; complete implementation details settings can be found in appendices B  
 and C. To ensure a fair comparison against existing methods, we trained separate model for each

270  
271 Table 1: Performance comparison on unconditional 3D molecular generation. Results for UniGEM are marked  
with an asterisk (\*) to indicate the use of additional molecular property information during training.

272 273	274	QM9				DRUG	
		Atom Sta(%))↑	Mol Sta(%))↑	Valid(%))↑	V × U(%))↑	Atom Sta(%))↑	Valid(%))↑
	Data	99.0	95.2	97.7	97.7	86.5	99.9
275	ENF (Garcia Satorras et al., 2021)	85.0	4.9	40.2	39.4	-	-
276	G-Schnet (Gebauer et al., 2022)	95.7	68.1	85.5	80.3	-	-
277	GDM (Hoogeboom et al., 2022)	97.0	63.2	-	-	75.0	90.8
278	GDM-AUG (Hoogeboom et al., 2022)	97.6	71.6	90.4	89.5	77.7	91.8
279	EDM (Hoogeboom et al., 2022)	98.7	82.0	91.9	90.7	81.3	92.6
280	EDM-Bridge (Wu et al., 2022)	98.8	84.6	92.0	90.7	82.4	92.8
	GeoLDM (Xu et al., 2023b)	98.9	89.4	93.8	92.7	84.4	99.3
	UniGEM* (Feng et al., 2024)	99.0	89.8	95.0	93.2	85.1	98.4
281	Uni-3DAR	<b>99.4</b>	<b>93.7</b>	<b>98.0</b>	<b>94.0</b>	<b>85.5</b>	<b>99.4</b>

282 Table 2: Results on de novo crystal generation. Baseline results are taken from Xie et al. (2021).

283 284	285	Data	Method	Validity (%)↑		Coverage (%)↑		Property ↓		
				Struc.	Comp.	COV-R	COV-P	$d_p$	$d_E$	$d_{\text{clem}}$
286 287 288 289 290	Carbon-24	FTCP (Ren et al., 2021)	0.08	-	0.00	0.00	5.206	19.05	-	-
		G-SchNet (Gebauer et al., 2019)	99.94	-	0.00	0.00	0.9427	1.320	-	-
		P-G-SchNet (Gebauer et al., 2019)	48.39	-	0.00	0.00	1.533	134.7	-	-
		CDVAE (Xie et al., 2021)	<b>100.0</b>	-	99.80	83.08	0.1407	0.2850	-	-
		DiffCSP (Jiao et al., 2023)	<b>100.0</b>	-	<u>99.90</u>	<u>97.27</u>	<u>0.0805</u>	0.0820	-	-
		Uni-3DAR	99.99	-	<b>100.0</b>	<b>98.16</b>	<b>0.0660</b>	<b>0.0289</b>	-	-
291 292 293 294 295	MP-20	FTCP (Ren et al., 2021)	1.55	48.37	4.72	0.09	23.71	160.9	0.7363	-
		G-SchNet (Gebauer et al., 2019)	99.65	75.96	38.33	99.57	3.034	42.09	0.6411	-
		P-G-SchNet (Gebauer et al., 2019)	77.51	76.40	41.93	99.74	4.04	2.448	0.6234	-
		CDVAE (Xie et al., 2021)	<b>100.0</b>	<u>86.70</u>	99.15	99.49	0.6875	0.2778	1.432	-
		DiffCSP (Jiao et al., 2023)	<b>100.0</b>	83.25	<b>99.71</b>	<u>99.76</u>	<u>0.3502</u>	<u>0.1247</u>	0.3398	-
		FlowMM (Miller et al., 2024)	96.85	83.19	99.49	99.58	<b>0.239</b>	-	0.083	-
		Uni-3DAR	99.89	<b>90.31</b>	<u>99.62</u>	<b>99.83</b>	0.4768	<b>0.1237</b>	<b>0.0694</b>	-

297 benchmark. We defer the investigation of joint training to future work. Uni-3DAR is robust to  
298 hyper-parameters, requiring no significant tuning and using a consistent setting across all tasks.

### 300 3.1 3D SMALL MOLECULE GENERATION

301 We assess Uni-3DAR on unconditional 3D molecular generation, a fundamental task challenged by  
302 the need to produce realistic conformations while accounting for molecular flexibility and diverse  
303 rotatable bonds. Our evaluation employs two standard benchmarks: **QM9** (Ramakrishnan et al.,  
304 2014b), a dataset of small molecules with up to 29 atoms, and **GEOM-DRUG** (Axelrod and Gomez-  
305 Bombarelli, 2022), which contains larger, more complex drug-like compounds with up to 181 atoms.  
306 Following the established protocols of Hoogeboom et al. (2022), we report on key metrics including  
307 Atom Stability, Molecule Stability, chemical validity (as determined by RDKit), and uniqueness.  
308 Bond types are inferred from the generated geometries to evaluate chemical correctness.

309 As shown in table 1, Uni-3DAR significantly outperforms all baseline models. On QM9, it achieves  
310 notable improvements in crucial metrics, reaching a Molecule Stability of 93.7% and a Validity  
311 of 98.0%, substantially exceeding the second-best methods. These results underscore Uni-3DAR’s  
312 robust capability to generate high-quality, chemically valid molecules. Furthermore, Uni-3DAR sur-  
313 passes UniGEM, a model that leverages additional molecular property information during training,  
314 using only 3D geometric data. This highlights the efficacy and robustness of our proposed model.

### 315 3.2 CRYSTAL GENERATION

316 We evaluate Uni-3DAR on crystal structure generation, a task distinct from organic molecules due  
317 to crystals’ rigidity, symmetry, and periodicity. A crystal is represented by its lattice (parallelepiped  
318 unit cell) and atomic configurations. Uni-3DAR adopts a two-frame generation approach: first gen-  
319 erating lattice vertices, then atom positions within the lattice. We consider three tasks: (1) de novo  
320 crystal generation (unconditional sampling), (2) crystal structure prediction (CSP) from given com-  
321 positions, and (3) PXRD-guided CSP, which reconstructs crystal structures from PXRD signals and  
322 compositions, with practical relevance for real-world material discovery. For composition condition-  
323 ing, we prepend a token from a multi-hot atom-type vector. PXRD data ( $0^\circ$ – $120^\circ$  at  $0.1^\circ$  resolution)  
is converted into a 1200-dim vector, split into four segments, each as a conditional token—yielding

324 Table 3: Results on crystal structure prediction (CSP) and PXRD-guided CSP. For a fair comparison, we report  
 325 UniGenX results obtained from the model trained from scratch, rather than using its default configuration that  
 326 leverages large-scale datasets for additional pretraining and fine-tuning.

327 Method	328 Carbon-24		329 MPTS-52		330 MP-20		331 MP-20 (PXRD-Guided)	
	332 Match Rate (%) ↑	333 RMSE ↓	334 Match Rate (%) ↑	335 RMSE ↓	336 Match Rate (%) ↑	337 RMSE ↓	338 Match Rate (%) ↑	339 RMSE ↓
330 CDVAE (Xie et al., 2021)	17.09	0.2969	331 5.34	0.2106	332 33.90	0.1045	333 –	334 –
331 DiffCSP (Jiao et al., 2023)	17.54	0.2759	332 12.19	0.1786	333 51.49	0.0631	334 –	335 –
332 FlowMM (Miller et al., 2024)	23.47	0.4122	333 17.54	0.1726	334 61.39	0.0566	335 –	336 –
333 UniGenX (Zhang et al., 2025)	27.09	0.2264	334 29.09	0.1256	335 63.88	0.0598	336 –	337 –
334 PXRDGEN (Li et al., 2024a)	–	–	335 –	336 –	337 –	338 –	339 68.68	340 0.0707
339 Uni-3DAR	<b>31.23</b>	<b>0.2194</b>	340 <b>32.44</b>	341 <b>0.0684</b>	342 <b>65.48</b>	343 <b>0.0317</b>	344 <b>75.08</b>	345 <b>0.0276</b>

346 five tokens in total (one for composition, four for PXRD). Uni-3DAR’s autoregressive framework  
 347 integrates these tokens directly, avoiding extra encoders used in prior work (Li et al., 2024a; Lai  
 348 et al., 2025). Following prior work (Xie et al., 2021; Jiao et al., 2023; Miller et al., 2024), we use  
 349 Carbon-24 (Pickard, 2020), MP-20 (Jain et al., 2013), and MPTS-52 datasets. De novo generation  
 350 is evaluated via validity, coverage, and property statistics (Xie et al., 2021), while CSP and PXRD-  
 351 guided CSP are assessed by top-1 match rate and RMSE, using StructureMatcher (Ong et al.,  
 352 2013) with the same thresholds as in (Miller et al., 2024).

353 Table 2 shows Uni-3DAR’s performance on Carbon-24 and MP-20. On Carbon-24, Uni-3DAR  
 354 outperforms baselines, especially in coverage, generating diverse and realistic structures. On MP-20,  
 355 it achieves higher component validity while maintaining competitive results overall, highlighting its  
 356 strength in producing chemically valid crystals. Table 3 summarizes CSP results across all datasets.  
 357 Uni-3DAR consistently outperforms baselines, improving match rate by 4.14% on Carbon-24 and  
 358 reducing RMSE from 0.0566 to 0.0317 on MP-20 (178% relative gain). On MPTS-52, it achieves  
 359 0.0684 RMSE, a 184% improvement despite higher complexity, demonstrating strong precision and  
 360 generalization. For PXRD-guided CSP, Uni-3DAR surpasses PXRDGEN (Li et al., 2024a), raising  
 361 the match rate from 68.68% to 75.08% and cutting RMSE from 0.0707 to 0.0276 (256% relative  
 362 gain), showing exceptional accuracy in reconstructing crystals from PXRD data.

### 3.3 MACROSCOPIC 3D OBJECT GENERATION

363 To demonstrate its versatility beyond micro-  
 364scopic domains, Uni-3DAR was also eval-  
 365uated on unconditional macroscopic 3D object  
 366 generation, a fundamental task in 3D com-  
 367 puter vision. We utilized three ShapeNet cat-  
 368 egories (*airplane*, *chair*, *car*) (Chang et al.,  
 369 2015), where objects are represented as point  
 370 clouds, and assessed using 1-NNA (with both  
 371 Chamfer distance (CD) and earth mover dis-  
 372 tance (EMD) as our main metric following Vah-  
 373 dat et al. (2022). A distinctive aspect for this  
 374 task is Uni-3DAR’s processing of an input object as  $512 \times 512 \times 512$  voxels, and the resulting 3D  
 375 patches (fine-grained structural tokens) are defined as  $16 \times 16 \times 16$  voxels. Each patch is quan-  
 376 tized using VQVAE. As shown in Table 4, Uni-3DAR exhibits highly competitive, often superior,  
 377 performance against established baselines (Yang et al., 2019). More details are in Appendix B.

### 3.4 PROTEIN POCKET PREDICTION

378 Predicting protein binding pockets is crucial for drug design and molecular docking. We eval-  
 379 uate Uni-3DAR’s token-level understanding on this task, formulating it as a classical atom-level  
 380 classification problem where each atom is labeled as part of a pocket or not. Following previous  
 381 work (Zhao et al., 2024), we train and evaluate on a composite dataset built from the CASF-2016  
 382 core set (Su et al., 2018), the PDBBind v2020 refined set (pdb, 2025), and MOAD (Hu et al., 2005).  
 383 Performance is measured using the Intersection-over-Union (IoU) metric. As shown in Table 6, Uni-  
 384 3DAR achieves state-of-the-art performance. Notably, it matches or exceeds specialized methods  
 385 like Vabs-Net, which relies on additional features such as ESM embeddings and Solvent Accessible  
 386 Surface Area, whereas Uni-3DAR uses only 3D structural information. These results highlight Uni-  
 387 3DAR’s strong capacity to interpret protein structures for fine-grained, atom-level prediction tasks.

388 Table 4: Unconditional 3D object generation results (1-  
 389 NNA↓) on ShapeNet. The **best** and second-best results  
 390 among the baselines are highlighted.

Method	Airplane		Chair		Car	
	CD ↓	EMD ↓	CD ↓	EMD ↓	CD ↓	EMD ↓
r-GAN (Achlioptas et al., 2018)	98.40	96.79	83.69	99.70	94.46	99.01
I-GAN (CD) (Achlioptas et al., 2018)	87.30	93.95	68.58	83.84	66.49	88.78
I-GAN (EMD) (Achlioptas et al., 2018)	89.49	76.91	71.90	64.65	71.16	66.19
PointFlow (Yang et al., 2019)	75.68	70.74	62.84	60.57	58.10	56.25
SoftFlow (Kim et al., 2020)	76.05	65.80	59.21	60.05	64.77	60.09
SetVAE (Kim et al., 2021)	76.54	67.65	58.84	60.57	59.94	59.94
DPF-Net (Kliokov et al., 2020)	75.18	65.55	62.00	58.53	62.35	54.48
DPM (Luo and Hu, 2021)	76.42	86.91	60.05	74.77	68.89	79.97
PVD (Zhou et al., 2021)	73.82	64.81	56.26	53.32	54.55	53.83
LION (Vahdat et al., 2022)	67.41	61.23	53.70	52.34	53.41	51.14
Uni-3DAR (Ours)	<b>67.35</b>	<b>61.09</b>	<b>53.11</b>	<b>51.78</b>	<b>53.35</b>	<b>50.89</b>

378 3.5 MOLECULAR DOCKING  
379

380 Molecular docking, which predicts the binding pose of a ligand to a protein, is a cornerstone of  
381 drug discovery. Uni-3DAR frames this as a three-frame generation task: the first two frames are the  
382 protein and the initial ligand conformation, and the third is the predicted docked pose. We evaluate  
383 this approach on the PDBbind2020 dataset (pdb, 2025), benchmarking against 13 classical and deep  
384 learning methods using standard RMSD-based metrics (Top-1/5 success rates for  $\text{RMSD} < 1\text{\AA}$  and  
385  $< 2\text{\AA}$ , and median RMSD), following the protocol of Cao et al. (2024). Uni-3DAR operates solely  
386 on atom types and coordinates, forgoing complex feature engineering and a separate scoring model;  
387 poses are ranked using the cumulative probability from the autoregressive generation. The results in  
388 Table 7 demonstrate that Uni-3DAR achieves state-of-the-art performance. It surpasses the previous  
389 best, SurfDock, on Top-1 metrics, with higher success rates for poses with  $\text{RMSD} < 1\text{\AA}$  (44.75%  
390 vs. 40.96%) and  $< 2\text{\AA}$  (69.06% vs. 68.41%), and a lower median RMSD ( $1.08\text{\AA}$  vs.  $1.18\text{\AA}$ ). While  
391 its Top-5 performance is slightly lower, likely due to its implicit scoring mechanism, these results  
392 underscore the strong potential of our unified, feature-light approach for molecular docking.

393 3.6 MOLECULAR AND POLYMER PROPERTY PREDICTION VIA PRETRAINING  
394

395 To evaluate its structure-level understanding, we assess Uni-3DAR on property prediction for small  
396 molecules and homopolymers after pretraining. For small molecules, we adopt the pretraining data,  
397 downstream tasks, and evaluation settings from state-of-the-art models Uni-Mol (Zhou et al., 2023b)  
398 and SpaceFormer (Lu et al., 2025), using Mean Absolute Error (MAE) as the metric. For homopolymers,  
399 we follow Wang et al. (2024) and use eight DFT-calculated property datasets, reporting the  
400 Root Mean Squared Error (RMSE) from a 5-fold cross-validation averaged over three seeds.

401 As summarized in Tables 8 and 9, Uni-3DAR demonstrates strong and versatile performance.  
402 On small molecule tasks (Table 8), it ranks first in 4 of 10 tasks and in the top two for 8 of 10,  
403 performing comparably to the specialized SpaceFormer model. On homopolymer tasks (Table 9), it  
404 ranks first in 4 of 8 tasks and in the top two for 7 of 8. These results affirm that Uni-3DAR develops  
405 robust and competitive representations for predicting properties across diverse chemical systems.

406 3.7 ADDITIONAL EXPERIMENTS  
407

408 Due to space limitations, we present further experimental results in Appendix D. These include (1)  
409 an analysis of the benefits of unifying understanding and generation, (2) comprehensive ablation  
410 studies evaluating our proposed tokenization and MNTP, and (3) a comparison of inference speeds.

411 4 RELATED WORK  
412

413 **Octree and Hierarchical Autoregressive Models** The coarse-to-fine hierarchical structure is  
414 widely used in 3D vision (Wang et al., 2017; Tatarchenko et al., 2017; Tang et al., 2021; Zhou et al.,  
415 2023a; Wang, 2023; Ibing et al., 2023; Zhang et al., 2024b; Ren et al., 2024). Among these works,  
416 (Ibing et al., 2023) is most similar to Uni-3DAR, as it also employs autoregressive generation using  
417 an octree. However, our method differs in three key aspects: (1) instead of relying on deep tree-  
418 level generation for fine details, we add an extra layer of fine-grained tokens to avoid excessively  
419 deep trees; (2) rather than compressing nodes via convolutional layers, we represent a compressed  
420 subtree with a single token; and (3) to handle dynamic token positions, while (Ibing et al., 2023)  
421 appends the next position to the current token, we adopt a masked next-token prediction strategy.  
422 These innovations make Uni-3DAR more efficient and effective than (Ibing et al., 2023). Recently,  
423 some image generative models have adopted a coarse-to-fine, level-by-level generation approach,  
424 such as VAR (Tian et al., 2024). Although the high-level idea appears similar, our motivation is dis-  
425 tinct: Uni-3DAR is designed to avoid the inefficiencies of a full-size cubic grid, whereas VAR uses  
426 more tokens to boost performance. Moreover, Uni-3DAR remains within the next-token prediction  
427 framework, while VAR employs next-scale prediction.

428 **Microscopic 3D Structure Modeling** Most previous generative models for microscopic 3D  
429 structures employ diffusion-based approaches (Wu et al., 2022; Anand and Achim, 2022; Hooge-  
430 boom et al., 2022; Xu et al., 2023b; Jiao et al., 2023) to generate atomic positions from noise.  
431 However, diffusion models have two major limitations. First, they require the number of atoms  
432 to be predetermined. Second, atom types are sampled from a categorical distribution, for which a  
433 proper score function is not well defined. Some studies have explored grid-based generation (O Pin-  
434 heiro et al., 2024), but using a full-size 3D grid is computationally prohibitive. Other works have

432 investigated autoregressive models for 3D molecules (Luo and Ji, 2022; Luo et al., 2021; Zhang  
 433 et al., 2025), but these models generate molecules atom by atom, requiring a predefined sequence  
 434 order. For microscopic 3D structure understanding, prior studies primarily leverage SE(3)-invariant  
 435 or equivariant models (Schütt et al., 2021; Fuchs et al., 2020). Additionally, unsupervised pretrain-  
 436 ing is widely used to mitigate the scarcity of labeled data (Stärk et al., 2022a; Cui et al., 2024; Yang  
 437 et al., 2024; Zaidi et al., 2022; Zhou et al., 2023b). These models typically follow a BERT-style pre-  
 438 training framework (Devlin, 2018), where some atoms are masked, their 3D positions are perturbed,  
 439 and the model is trained to recover the ground truth. While highly effective for understanding tasks,  
 440 most of these models cannot be directly applied to generation.

441 Some recent efforts have attempted to unify generation and understanding for microscopic data.  
 442 However, most focus solely on sequence data (e.g., 1D SMILES, nucleotide sequences, or textual  
 443 descriptions) and directly apply autoregressive language models (Christofidellis et al., 2023; Zhang  
 444 et al., 2024a; Nguyen et al., 2024; Xia et al., 2025). While these models are straightforward, they  
 445 lack essential 3D structural information, limiting their performance and applicability. Recent studies  
 446 have also explored diffusion-based approaches. For example, UniGEM (Feng et al., 2024) demon-  
 447 strated that a two-phase, multi-task training strategy can improve performance for both tasks. This  
 448 approach combines diffusion loss with a prediction task applied during later diffusion steps. In sum-  
 449 mary, while previous work has made progress in bridging generation and understanding, Uni-3DAR  
 450 is the first autoregressive framework to unify both tasks for 3D microscopic structures.

451 **Macroscopic 3D Structure Modeling** Macroscopic 3D structure modeling encompasses the  
 452 understanding and generation of everyday objects (Chang et al., 2015; Deitke et al., 2023),  
 453 scenes (Peng et al., 2023), CAD models (Wu et al., 2021; Willis et al., 2021; Xu et al., 2024),  
 454 avatars (Canfes et al., 2023), and more. Similar to microscopic 3D structures, macroscopic 3D  
 455 structures lack a unified representation format. Commonly used 3D representations include vox-  
 456 els (Wang et al., 2018b), point clouds (Xue et al., 2023), polygon meshes (Liu et al., 2024), implicit  
 457 functions (Tang et al., 2021), and 3D Gaussian Splatting (Kerbl et al., 2023). Recent methods (Zhang  
 458 et al., 2023a; Zhao et al., 2023; Zhang et al., 2024b; Chen et al., 2024) based on Diffusion Trans-  
 459 formers (Peebles and Xie, 2023) encode 3D shapes into compressed, compact latent codes, substantially  
 460 improving representation efficiency. Previous literature also explored autoregressive modeling for  
 461 macroscopic 3D structures. For example, Polygen (Nash et al., 2020) and MeshGPT (Siddiqui et al.,  
 462 2024) generate mesh faces sequentially from lowest to highest on the vertical axis, correspond-  
 463 ing to the point-based tokenization strategy as discussed in Sec. 1, suffering from the same chal-  
 464 lenges in dynamic token positions. Another category of 3D structure generation methods, known as  
 465 optimization-based approaches (Tang et al., 2023; Lin et al., 2023a; Metzer et al., 2023; Poole et al.,  
 466 2022), leverages text-to-image generative models and refines 3D representations by distilling infor-  
 467 mation from 2D images (Poole et al., 2022). Unlike true 3D generation, these methods primarily  
 468 perform 3D reconstruction, making them fundamentally distinct from the previously mentioned 3D  
 469 generation techniques and Uni-3DAR.

## 470 5 CONCLUSION

471 In this work, we introduced Uni-3DAR, a unified autoregressive framework designed to address the  
 472 long-standing fragmentation of 3D modeling. By leveraging a novel coarse-to-fine octree-based tok-  
 473 enizer, Uni-3DAR compresses diverse 3D structures—from molecules to macroscopic shapes—into  
 474 a common 1D sequence representation. This core innovation, enhanced by 2-level subtree compres-  
 475 sion for efficiency and a masked next-token prediction strategy to handle sparse spatial data, enables  
 476 a single model to seamlessly bridge the gap between generative and understanding tasks across  
 477 different scales. Our extensive experiments validate this unified approach, demonstrating that Uni-  
 478 3DAR achieves state-of-the-art or highly competitive performance on a wide array of benchmarks.  
 479 Notably, it consistently outperforms specialized, diffusion-based models while being significantly  
 480 more efficient. **Additionally, we prioritize a general, scalable architecture over hard-coded inductive  
 481 biases like SE(3) invariance. By relying on data augmentation rather than architectural constraints,  
 482 Uni-3DAR effectively learns geometric principles from data, matching or outperforming specialized  
 483 SE(3)-equivariant baselines on molecular tasks.** By proving that a simple yet powerful autoregres-  
 484 sive paradigm can unify disparate tasks without compromising accuracy, we believe Uni-3DAR  
 485 marks a pivotal step toward a general-purpose foundation model (a “GPT-2 moment”) for the cross-  
 scale 3D domain.

486     **Limitations** While our results demonstrate the mutual benefits of unifying generation and under-  
 487     standing, we have not yet trained a single, large-scale foundation model on a heterogeneous mixture  
 488     of 3D data and tasks. Realizing this vision through joint pretraining is a primary goal for future work.  
 489     Other critical avenues for research include extending the framework to real-world applications.  
 490

491     **ETHICS STATEMENT**  
 492

493     The research presented in this paper aims to advance scientific discovery by creating a unified frame-  
 494     work for 3D modeling. We have strived to conduct this work with the highest ethical standards.  
 495

496     All datasets used in our experiments—including QM9, GEOM-DRUG, Materials Project, ShapeNet,  
 497     and PDBBind—are publicly available and are standard benchmarks in their respective scientific  
 498     communities. We did not collect any new data, and no personally identifiable or sensitive informa-  
 499     tion was used.

500     We acknowledge that generative models for molecular and material design could potentially be mis-  
 501     used for creating harmful substances. However, Uni-3DAR is intended as a fundamental scientific  
 502     tool to accelerate beneficial research in fields such as drug discovery and materials science. Its ca-  
 503     pabilities are grounded in the principles learned from public scientific data. We believe the potential  
 504     benefits—such as the rapid design of novel therapeutics and efficient materials—significantly out-  
 505     weigh the risks. As with any powerful technology, we advocate for its responsible use and encourage  
 506     the research community to establish clear guidelines for the ethical application of generative models  
 507     in science.

508     Finally, we recognize the environmental impact associated with training large-scale models. Our  
 509     work incorporates significant efficiency optimizations, such as octree and subtree compression,  
 510     which substantially reduce the computational resources and token count required for training and  
 511     inference compared to alternative approaches.

512  
 513     **REPRODUCIBILITY STATEMENT**  
 514

515     To ensure the reproducibility of our results, we are committed to making our research as transparent  
 516     as possible.  
 517

518     **Code** Upon acceptance of this paper, we will release the complete source code for Uni-3DAR,  
 519     including model implementation, training scripts, and evaluation protocols, under a permissive open-  
 520     source license.  
 521

522     **Data** All datasets used in our experiments are publicly available and have been cited appropriately.  
 523     We followed standard data processing and splitting protocols as established in prior work. Detailed  
 524     descriptions of data preparation for each task are provided in Appendix C.  
 525

526     **Hyperparameters and Architecture** The full details of our model architecture, as well as the  
 527     specific hyperparameters used for every experiment (including learning rates, batch sizes, model  
 528     dimensions, and training steps), are thoroughly documented in Appendix B and C. We used a con-  
 529     sistent model configuration across most tasks to demonstrate the robustness and generality of our  
 530     framework.  
 531

532     **Computational Environment** Our experiments were conducted using standard deep learning li-  
 533     braries. Specific details about the hardware (e.g., NVIDIA A100 and 4090 GPUs) and software  
 534     environment are provided in the appendices to facilitate the replication of our training and infer-  
 535     ence setups. The efficiency optimizations used, such as FlashAttention and KV-caching, are also  
 536     described in Appendix B.  
 537

538     **REFERENCES**  
 539

Pdbbind+, 2025. URL <https://www.pdbbind-plus.org.cn/>. 7, 8, 26

540 Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations  
 541 and generative models for 3d point clouds. In *International conference on machine learning*,  
 542 pages 40–49. PMLR, 2018. 7, 25

543

544 Eric Alcaide, Zhifeng Gao, Guolin Ke, Yaqi Li, Linfeng Zhang, Hang Zheng, and Gengmo Zhou.  
 545 Uni-mol docking v2: Towards realistic and accurate binding pose prediction. *arXiv preprint*  
 546 *arXiv:2405.11769*, 2024. 1

547

548 Dosovitskiy Alexey. An image is worth 16x16 words: Transformers for image recognition at scale.  
 549 *arXiv preprint arXiv: 2010.11929*, 2020. 4

550

551 Namrata Anand and Tudor Achim. Protein structure and sequence generation with equivariant de-  
 552 noising diffusion probabilistic models. *arXiv preprint arXiv:2205.15019*, 2022. 8

553

554 Simon Axelrod and Rafael Gomez-Bombarelli. Geom, energy-annotated molecular conformations  
 555 for property prediction and molecular generation. *Scientific Data*, 9(1):185, 2022. 6, 23

556

557 Zehranaz Canfes, M Furkan Atasoy, Alara Dirik, and Pinar Yanardag. Text and image guided 3d  
 558 avatar generation and manipulation. In *Proceedings of the IEEE/CVF Winter Conference on*  
 559 *Applications of Computer Vision*, pages 4421–4431, 2023. 9

560

561 Duanhua Cao, Mingan Chen, Runze Zhang, Zhaokun Wang, Manlin Huang, Jie Yu, Xinyu Jiang,  
 562 Zhehuan Fan, Wei Zhang, Hao Zhou, et al. Surfdock is a surface-informed diffusion generative  
 563 model for reliable and accurate protein–ligand complex prediction. *Nature Methods*, pages 1–13,  
 564 2024. 8, 27

565

566 Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li,  
 567 Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d  
 568 model repository. *arXiv preprint arXiv:1512.03012*, 2015. 7, 9, 25

569

570 Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative  
 571 image transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
 572 *recognition*, pages 11315–11325, 2022. 5

573

574 Rui Chen, Jianfeng Zhang, Yixun Liang, Guan Luo, Weiyu Li, Jiarui Liu, Xiu Li, Xiaoxiao Long,  
 575 Jiashi Feng, and Ping Tan. Dora: Sampling and benchmarking for 3d shape variational auto-  
 576 encoders. *arXiv preprint arXiv:2412.17808*, 2024. 9

577

578 Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta: Large-scale self-  
 579 supervised pretraining for molecular property prediction, 2020. URL <https://arxiv.org/abs/2010.09885>. 29

580

581 Dimitrios Christofidellis, Giorgio Giannone, Jannis Born, Ole Winther, Teodoro Laino, and Matteo  
 582 Manica. Unifying molecular and textual representations via multi-task language modelling. In  
 583 *International Conference on Machine Learning*, pages 6140–6157. PMLR, 2023. 9

584

585 Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi S Jaakkola. Diffdock:  
 586 Diffusion steps, twists, and turns for molecular docking. In *The Eleventh International Conference*  
 587 *on Learning Representations*. 27

588

589 Gabriele Corso, Arthur Deng, Benjamin Fry, Nicholas Polizzi, Regina Barzilay, and Tommi  
 590 Jaakkola. Deep confident steps to new pockets: Strategies for docking generalization. *ArXiv*,  
 591 pages arXiv–2402, 2024. 27

592

593 Taoyong Cui, Chenyu Tang, Mao Su, Shufei Zhang, Yuqiang Li, Lei Bai, Yuhang Dong, Xingao  
 594 Gong, and Wanli Ouyang. Geometry-enhanced pretraining on interatomic potentials. *Nature*  
 595 *Machine Intelligence*, 6(4):428–436, 2024. 9

596

597 Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and  
 598 memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing*  
 599 *Systems (NeurIPS)*, 2022. 22

594 Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig  
 595 Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-  
 596 tated 3d objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
 597 *recognition*, pages 13142–13153, 2023. 9

598  
 599 Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding.  
 600 *arXiv preprint arXiv:1810.04805*, 2018. 9

601 Jerome Eberhardt, Diogo Santos-Martins, Andreas F Tillack, and Stefano Forli. Autodock vina  
 602 1.2. 0: New docking methods, expanded force field, and python bindings. *Journal of chemical*  
 603 *information and modeling*, 61(8):3891–3898, 2021. 27

604  
 605 Xiaomin Fang, Lihang Liu, Jieqiong Lei, Donglong He, Shanzhuo Zhang, Jingbo Zhou, Fan Wang,  
 606 Hua Wu, and Haifeng Wang. Geometry-enhanced molecular representation learning for property  
 607 prediction. *Nature Machine Intelligence*, 4(2):127–134, 2022. 28

608 Shikun Feng, Yuyan Ni, Yan Lu, Zhi-Ming Ma, Wei-Ying Ma, and Yanyan Lan. Unigem: A unified  
 609 approach to generation and property prediction for molecules. *arXiv preprint arXiv:2410.10516*,  
 610 2024. 6, 9, 23

611  
 612 Richard A Friesner, Jay L Banks, Robert B Murphy, Thomas A Halgren, Jasna J Klicic, Daniel T  
 613 Mainz, Matthew P Repasky, Eric H Knoll, Mee Shelley, Jason K Perry, et al. Glide: a new  
 614 approach for rapid, accurate docking and scoring. 1. method and assessment of docking accuracy.  
 615 *Journal of medicinal chemistry*, 47(7):1739–1749, 2004. 27

616 Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-  
 617 translation equivariant attention networks. *Advances in neural information processing systems*,  
 618 33:1970–1981, 2020. 9

619  
 620 Victor Garcia Satorras, Emiel Hoogeboom, Fabian Fuchs, Ingmar Posner, and Max Welling. E (n)  
 621 equivariant normalizing flows. *Advances in Neural Information Processing Systems*, 34:4181–  
 622 4192, 2021. 6, 23

623 Niklas Gebauer, Michael Gastegger, and Kristof Schütt. Symmetry-adapted generation of 3d point  
 624 sets for the targeted discovery of molecules. In H. Wallach, H. Larochelle, A. Beygelzimer,  
 625 F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Sys-*  
 626 *tems* 32, pages 7566–7578. Curran Associates, Inc., 2019. 6, 24

627  
 628 Niklas WA Gebauer, Michael Gastegger, Stefaan SP Hessmann, Klaus-Robert Müller, and Kristof T  
 629 Schütt. Inverse design of 3d molecular structures with conditional generative neural networks.  
 630 *Nature communications*, 13(1):973, 2022. 6, 23

631 Marylens Hernandez, Dario Ghersi, and Roberto Sanchez. Sitehound-web: a server for ligand bind-  
 632 ing site identification in protein structures. *Nucleic acids research*, 37(suppl\_2):W413–W416,  
 633 2009. 26

634  
 635 Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant dif-  
 636 fusion for molecule generation in 3d. In *International conference on machine learning*, pages  
 637 8867–8887. PMLR, 2022. 6, 8, 23

638 Liegi Hu, Mark L Benson, Richard D Smith, Michael G Lerner, and Heather A Carlson. Binding  
 639 moad (mother of all databases). *Proteins: Structure, Function, and Bioinformatics*, 60(3):333–  
 640 340, 2005. 7, 26

641  
 642 Moritz Ibing, Gregor Kobsik, and Leif Kobbelt. Octree transformer: Autoregressive 3d shape gen-  
 643 eration on hierarchically structured sequences. In *Proceedings of the IEEE/CVF Conference on*  
 644 *Computer Vision and Pattern Recognition*, pages 2698–2707, 2023. 4, 5, 8, 30

645 Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen  
 646 Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The ma-  
 647 terials project: A materials genome approach to accelerating materials innovation. *APL materials*,  
 1(1):011002, 2013. 7, 24

648 Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal struc-  
 649 ture prediction by joint equivariant diffusion on lattices and fractional coordinates. In *Workshop*  
 650 on "Machine Learning for Materials" ICLR 2023, 2023. URL [https://openreview.net/](https://openreview.net/forum?id=VPByphdu24j)  
 651 [forum?id=VPByphdu24j](https://openreview.net/forum?id=VPByphdu24j). 1, 6, 7, 8, 24

652 José Jiménez, Stefan Doerr, Gerard Martínez-Rosell, Alexander S Rose, and Gianni De Fabritiis.  
 653 Deepsite: protein-binding site predictor using 3d-convolutional neural networks. *Bioinformatics*,  
 654 33(19):3036–3042, 2017. 26

655 John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger,  
 656 Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate  
 657 protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021. 1

658 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-  
 659 ting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 9

660 Hyeongju Kim, Hyeonseung Lee, Woo Hyun Kang, Joun Yeop Lee, and Nam Soo Kim. Softflow:  
 661 Probabilistic framework for normalizing flow on manifolds. *Advances in Neural Information  
 662 Processing Systems*, 33:16388–16397, 2020. 7, 25

663 Jinwoo Kim, Jaehoon Yoo, Juho Lee, and Seunghoon Hong. Setvae: Learning hierarchical compo-  
 664 sition for generative modeling of set-structured data. In *Proceedings of the IEEE/CVF Conference  
 665 on Computer Vision and Pattern Recognition*, pages 15059–15068, 2021. 7, 25

666 Roman Klokov, Edmond Boyer, and Jakob Verbeek. Discrete point flow networks for efficient point  
 667 cloud generation. In *European Conference on Computer Vision*, pages 694–710. Springer, 2020.  
 668 7, 25

669 David Ryan Koes, Matthew P Baumgartner, and Carlos J Camacho. Lessons learned in empirical  
 670 scoring with smina from the csar 2011 benchmarking exercise. *Journal of chemical information  
 671 and modeling*, 53(8):1893–1904, 2013. 27

672 Radoslav Krivák and David Hoksza. P2rank: machine learning based tool for rapid and accurate  
 673 prediction of ligand binding sites from protein structure. *Journal of cheminformatics*, 10:1–12,  
 674 2018. 26

675 Christopher Kuenneth and Rampi Ramprasad. polybert: a chemical language model to enable fully  
 676 machine-driven ultrafast polymer informatics. *Nature Communications*, 14(1):4099, 2023. 29

677 Qingsi Lai, Fanjie Xu, Lin Yao, Zhifeng Gao, Siyuan Liu, Hongshuai Wang, Lu Shuqi, Di He, Liwei  
 678 Wang, Linfeng Zhang, Cheng Wang, and Guolin Ke. End-to-end crystal structure prediction from  
 679 powder x-ray diffraction. *Advanced Science*, 12, 01 2025. doi: 10.1002/advs.202410722. 7, 24

680 Vincent Le Guilloux, Peter Schmidke, and Pierre Tuffery. Fpocket: An open source platform  
 681 for ligand pocket detection. *Bioinformatics*, 10(1):168, 2009. ISSN 1471-2105. doi: 10.1186/  
 682 1471-2105-10-168. 26

683 Qi Li, Rui Jiao, Liming Wu, Tiannian Zhu, Wenbing Huang, Shifeng Jin, Yang Liu, Hongming  
 684 Weng, and Xiaolong Chen. Powder diffraction crystal structure determination using generative  
 685 models, 2024a. URL <https://arxiv.org/abs/2409.04727>. 7, 24, 25

686 Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image  
 687 generation without vector quantization. *arXiv preprint arXiv:2406.11838*, 2024b. 4, 5, 22, 31

688 Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten  
 689 Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content  
 690 creation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,  
 691 pages 300–309, 2023a. 9

692 Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Allan dos San-  
 693 tos Costa, Maryam Fazel-Zarandi, Tom Sercu, Sal Candido, et al. Language models of protein  
 694 sequences at the scale of evolution enable accurate structure prediction. *BioRxiv*, 2022:500902,  
 695 2022. 27

702 Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin,  
 703 Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level  
 704 protein structure with a language model. *Science*, 379(6637):1123–1130, 2023b. 26

705 Minghua Liu, Chong Zeng, Xinyue Wei, Ruoxi Shi, Linghao Chen, Chao Xu, Mengqi Zhang,  
 706 Zhaoning Wang, Xiaoshuai Zhang, Isabella Liu, et al. Meshformer: High-quality mesh gener-  
 707 ation with 3d-guided reconstruction model. *arXiv preprint arXiv:2408.10198*, 2024. 9

708 Shuqi Lu, Xiaohong Ji, Bohang Zhang, Lin Yao, Siyuan Liu, Zhifeng Gao, Linfeng Zhang, and  
 709 Guolin Ke. Beyond atoms: Enhancing molecular pretrained representations with 3d space mod-  
 710eling. *arXiv preprint arXiv:2503.10489*, 2025. 8, 28

711 Wei Lu, Qifeng Wu, Jixian Zhang, Jiahua Rao, Chengtao Li, and Shuangjia Zheng. Tankbind:  
 712 Trigonometry-aware neural networks for drug-protein binding structure prediction. *Advances in  
 713 neural information processing systems*, 35:7236–7249, 2022. 27

714 Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceed-  
 715 ings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2837–2845,  
 716 2021. 7, 25

717 Shitong Luo, Jiaqi Guan, Jianzhu Ma, and Jian Peng. A 3d generative model for structure-based  
 718 drug design. *Advances in Neural Information Processing Systems*, 34:6229–6239, 2021. 9

719 Youzhi Luo and Shuiwang Ji. An autoregressive flow model for 3d molecular geometry generation  
 720 from scratch. In *International conference on learning representations (ICLR)*, 2022. 9

721 Gabriele Macari, Daniele Toti, and Fabio Polticelli. Computational methods and tools for binding  
 722 site recognition between proteins and small molecules: from classical geometrical approaches to  
 723 modern machine learning strategies. *Journal of computer-aided molecular design*, 33:887–903,  
 724 2019. 26

725 Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for  
 726 shape-guided generation of 3d shapes and textures. In *Proceedings of the IEEE/CVF conference  
 727 on computer vision and pattern recognition*, pages 12663–12673, 2023. 9

728 Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M Wood. Flowmm: Gen-  
 729 erating materials with riemannian flow matching, 2024. URL <https://arxiv.org/abs/2406.04713>. 6, 7, 24

730 Charlie Nash, Yaroslav Ganin, SM Ali Eslami, and Peter Battaglia. Polygen: An autoregressive  
 731 generative model of 3d meshes. In *International conference on machine learning*, pages 7220–  
 732 7229. PMLR, 2020. 9

733 Eric Nguyen, Michael Poli, Matthew G Durrant, Brian Kang, Dhruva Katrekar, David B Li, Liam J  
 734 Bartie, Armin W Thomas, Samuel H King, Garyk Bixi, et al. Sequence modeling and design  
 735 from molecular to genome scale with evo. *Science*, 386(6723):ead09336, 2024. 9

736 Pedro O O Pinheiro, Joshua Rackers, Joseph Kleinhenz, Michael Maser, Omar Mahmood, Andrew  
 737 Watkins, Stephen Ra, Vishnu Sresht, and Saeed Saremi. 3d molecule generation by denoising  
 738 voxel grids. *Advances in Neural Information Processing Systems*, 36:69077–69097, 2023. 2

739 Pedro O O Pinheiro, Joshua Rackers, Joseph Kleinhenz, Michael Maser, Omar Mahmood, Andrew  
 740 Watkins, Stephen Ra, Vishnu Sresht, and Saeed Saremi. 3d molecule generation by denoising  
 741 voxel grids. *Advances in Neural Information Processing Systems*, 36, 2024. 8

742 Shyue Ping Ong, William Davidson Richards, Anubhav Jain, Geoffroy Hautier, Michael Kocher,  
 743 Shreyas Cholia, Dan Gunter, Vincent L Chevrier, Kristin A Persson, and Gerbrand Ceder. Python  
 744 materials genomics (pymatgen): A robust, open-source python library for materials analysis.  
 745 *Computational Materials Science*, 68:314–319, 2013. 7, 24

746 Ziqi Pang, Tianyuan Zhang, Fujun Luan, Yunze Man, Hao Tan, Kai Zhang, William T Freeman, and  
 747 Yu-Xiong Wang. Randar: Decoder-only autoregressive visual generation in random orders. *arXiv  
 748 preprint arXiv:2412.01827*, 2024. 5

756 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*  
 757 *the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 9  
 758

759 Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas  
 760 Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. In *Proceed-  
 761 ings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 815–824,  
 762 2023. 9

763 Chris J. Pickard. Airss data for carbon at 10gpa and the c+n+h+o system at 1gpa, 2020. 7, 24  
 764

765 Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d  
 766 diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 9

767 Matthew Ragoza, Joshua Hochuli, Elisa Idrobo, Jocelyn Sunseri, and David Ryan Koes. Protein-  
 768 ligand scoring with convolutional neural networks. *Journal of chemical information and model-  
 769 ing*, 57(4):942–957, 2017. 27

770 Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum  
 771 chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014a. 21  
 772

773 Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum  
 774 chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014b. 6, 23  
 775

776 Xuanchi Ren, Jiahui Huang, Xiaohui Zeng, Ken Museth, Sanja Fidler, and Francis Williams.  
 777 Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies. In *Proceedings of the  
 778 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4209–4219, 2024. 8

779 Zekun Ren, Siyu Isaac Parker Tian, Juhwan Noh, Felipe Oviedo, Guangzong Xing, Jiali Li, Qiao-  
 780 hao Liang, Ruiming Zhu, Armin G. Aberle, Shijing Sun, Xiaonan Wang, Yi Liu, Qianxiao Li,  
 781 Senthilnath Jayavelu, Kedar Hippalgaonkar, Yousung Jung, and Tonio Buonassisi. An invert-  
 782 ible crystallographic representation for general inverse design of inorganic crystals with targeted  
 783 properties. *Matter*, 2021. ISSN 2590-2385. doi: <https://doi.org/10.1016/j.matt.2021.11.032>. 6,  
 784 24

785 Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang.  
 786 Self-supervised graph transformer on large-scale molecular data. *Advances in neural information  
 787 processing systems*, 33:12559–12571, 2020. 28

788

789 Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction  
 790 of tensorial properties and molecular spectra. In *International Conference on Machine Learning*,  
 791 pages 9377–9388. PMLR, 2021. 9

792 Noam Shazeer. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020. 22  
 793

794 Yawar Siddiqui, Antonio Alliegro, Alexey Artemov, Tatiana Tommasi, Daniele Sirigatti, Vladislav  
 795 Rosov, Angela Dai, and Matthias Nießner. Meshgpt: Generating triangle meshes with decoder-  
 796 only transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern  
 797 recognition*, pages 19615–19625, 2024. 9

798 Hannes Stärk, Dominique Beaini, Gabriele Corso, Prudencio Tossou, Christian Dallago, Stephan  
 799 Günemann, and Pietro Liò. 3d infomax improves gnns for molecular property prediction. In  
 800 *International Conference on Machine Learning*, pages 20479–20502. PMLR, 2022a. 9, 28  
 801

802 Hannes Stärk, Octavian Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola. Equibind:  
 803 Geometric deep learning for drug binding structure prediction. In *International conference on  
 804 machine learning*, pages 20503–20521. PMLR, 2022b. 27

805 Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: En-  
 806 hanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 22  
 807

808 Minyi Su, Qifan Yang, Yu Du, Guoqin Feng, Zhihai Liu, Yan Li, and Renxiao Wang. Comparative  
 809 assessment of scoring functions: the casf-2016 update. *Journal of chemical information and  
 modeling*, 59(2):895–913, 2018. 7, 26

810 Jia-Heng Tang, Weikai Chen, Jie Yang, Bo Wang, Songrun Liu, Bo Yang, and Lin Gao. Octfield:  
 811 Hierarchical implicit functions for 3d modeling. *arXiv preprint arXiv:2111.01067*, 2021. 8, 9  
 812

813 Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative  
 814 gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 9  
 815

816 Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient  
 817 convolutional architectures for high-resolution 3d outputs. In *Proceedings of the IEEE interna-*  
 818 *tional conference on computer vision*, pages 2088–2096, 2017. 8  
 819

820 Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling:  
 821 Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024. 8  
 822

823 Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, Karsten Kreis, et al. Lion: La-  
 824 tent point diffusion models for 3d shape generation. *Advances in Neural Information Processing  
 Systems*, 35:10021–10039, 2022. 7, 25  
 825

826 Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in  
 neural information processing systems*, 30, 2017. 1, 4  
 827

828 A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.  
 829 21  
 830

831 Fanmeng Wang, Wentao Guo, Minjie Cheng, Shen Yuan, Hongteng Xu, and Zhifeng Gao. Mm-  
 832 polymer: A multimodal multitask pretraining framework for polymer property prediction. In  
 833 *Proceedings of the 33rd ACM International Conference on Information and Knowledge Manage-  
 ment*, pages 2336–2346, 2024. 8, 29  
 834

835 Han Wang, Linfeng Zhang, Jiequn Han, et al. Deepmd-kit: A deep learning package for many-body  
 836 potential energy representation and molecular dynamics. *Computer Physics Communications*,  
 837 228:178–184, 2018a. 1  
 838

839 Peng-Shuai Wang. Octformer: Octree-based transformers for 3d point clouds. *ACM Transactions  
 on Graphics (TOG)*, 42(4):1–11, 2023. 8  
 840

841 Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based  
 842 convolutional neural networks for 3d shape analysis. *ACM Transactions On Graphics (TOG)*, 36  
 843 (4):1–11, 2017. 8  
 844

845 Peng-Shuai Wang, Chun-Yu Sun, Yang Liu, and Xin Tong. Adaptive o-cnn: A patch-based deep  
 846 representation of 3d shapes. *ACM Transactions on Graphics (TOG)*, 37(6):1–11, 2018b. 9  
 847

848 Karl DD Willis, Yewen Pu, Jieliang Luo, Hang Chu, Tao Du, Joseph G Lamourne, Armando Solar-  
 849 Lezama, and Wojciech Matusik. Fusion 360 gallery: A dataset and environment for programmatic  
 850 cad construction from human design sequences. *ACM Transactions on Graphics (TOG)*, 40(4):  
 1–24, 2021. 9  
 851

852 Lemeng Wu, Chengyue Gong, Xingchao Liu, Mao Ye, and Qiang Liu. Diffusion-based molecule  
 853 generation with informative prior bridges. *Advances in Neural Information Processing Systems*,  
 854 35:36533–36545, 2022. 6, 8, 23  
 855

856 Rundi Wu, Chang Xiao, and Changxi Zheng. Deepcad: A deep generative network for computer-  
 857 aided design models. In *Proceedings of the IEEE/CVF International Conference on Computer  
 Vision*, pages 6772–6782, 2021. 9  
 858

859 Yingce Xia, Peiran Jin, Shufang Xie, Liang He, Chuan Cao, Renqian Luo, Guoqing Liu, Yue Wang,  
 860 Zequn Liu, Yuan-Jyue Chen, et al. Naturelm: Deciphering the language of nature for scientific  
 861 discovery. *arXiv preprint arXiv:2502.07527*, 2025. 9  
 862

863 Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi S Jaakkola. Crystal  
 diffusion variational autoencoder for periodic material generation. In *International Conference  
 on Learning Representations*, 2021. 1, 6, 7, 24

864 Rubin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang,  
 865 Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture.  
 866 *In International conference on machine learning*, pages 10524–10533. PMLR, 2020. 22  
 867

868 Changwen Xu, Yuyang Wang, and Amir Barati Farimani. Transpolymer: a transformer-based lan-  
 869 guage model for polymer property predictions. *npj Computational Materials*, 9(1):64, 2023a.  
 870 29

871 Jingwei Xu, Chenyu Wang, Zibo Zhao, Wen Liu, Yi Ma, and Shenghua Gao. Cad-mllm: Unifying  
 872 multimodality-conditioned cad generation with mllm. *arXiv preprint arXiv:2411.04954*, 2024. 9  
 873

874 Minkai Xu, Alexander S Powers, Ron O Dror, Stefano Ermon, and Jure Leskovec. Geometric latent  
 875 diffusion models for 3d molecule generation. In *International Conference on Machine Learning*,  
 876 pages 38592–38610. PMLR, 2023b. 6, 8, 23, 31

877 Le Xue, Mingfei Gao, Chen Xing, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu,  
 878 Juan Carlos Niebles, and Silvio Savarese. Ulip: Learning a unified representation of language,  
 879 images, and point clouds for 3d understanding. In *Proceedings of the IEEE/CVF conference on*  
 880 *computer vision and pattern recognition*, pages 1179–1189, 2023. 9

881 Xinguang Yan, Liqiang Lin, Niloy J Mitra, Dani Lischinski, Daniel Cohen-Or, and Hui Huang.  
 882 Shapeformer: Transformer-based shape completion via sparse representation. In *Proceedings of*  
 883 *the IEEE/CVF conference on computer vision and pattern recognition*, pages 6239–6249, 2022.  
 884 5

885 Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan.  
 886 Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the*  
 887 *IEEE/CVF international conference on computer vision*, pages 4541–4550, 2019. 7, 25

888 Junwei Yang, Kangjie Zheng, Siyu Long, Zaiqing Nie, Ming Zhang, Xinyu Dai, Wei-Ying Ma, and  
 889 Hao Zhou. Mol-ae: Auto-encoder based molecular representation learning with 3d cloze test  
 890 objective. *bioRxiv*, pages 2024–04, 2024. 9, 28

891 Yuejiang Yu, Chun Cai, Zhengdan Zhu, and Hang Zheng. Uni-dock: A gpu-accelerated docking  
 892 program enables ultra-large virtual screening. 2022. 27

893 Sheheryar Zaidi, Michael Schaarschmidt, James Martens, Hyunjik Kim, Yee Whye Teh, Alvaro  
 894 Sanchez-Gonzalez, Peter Battaglia, Razvan Pascanu, and Jonathan Godwin. Pre-training via de-  
 895 noising for molecular property prediction. *arXiv preprint arXiv:2206.00133*, 2022. 9

896 Biao Zhang and Rico Sennrich. Root mean square layer normalization. *Advances in Neural Infor-*  
 897 *mation Processing Systems*, 32, 2019. 22

898 Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape  
 899 representation for neural fields and generative diffusion models. *ACM Transactions On Graphics*  
 900 (*TOG*), 42(4):1–16, 2023a. 9

901 Gongbo Zhang, Yanting Li, Renqian Luo, Pipi Hu, Zeru Zhao, Lingbo Li, Guoqing Liu, Zun Wang,  
 902 Ran Bi, Kaiyuan Gao, Liya Guo, Yu Xie, Chang Liu, Jia Zhang, Tian Xie, Robert Pinsler, Clau-  
 903 dio Zeni, Ziheng Lu, Yingce Xia, Marwin Segler, Maik Riechert, Li Yuan, Lei Chen, Haiguang  
 904 Liu, and Tao Qin. Unigenx: Unified generation of sequence and structure with autoregressive  
 905 diffusion, 2025. URL <https://arxiv.org/abs/2503.06687>. 7, 9, 24, 31

906 Juzheng Zhang, Yatao Bian, Yongqiang Chen, and Quanming Yao. Unimot: Unified molecule-text  
 907 language model with discrete token representation. *arXiv preprint arXiv:2408.00863*, 2024a. 9

908 Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan  
 909 Xu, and Jingyi Yu. Clay: A controllable large-scale generative model for creating high-quality 3d  
 910 assets. *ACM Transactions on Graphics (TOG)*, 43(4):1–20, 2024b. 8, 9

911 Pei Zhang, Logan Kearney, Debsindhu Bhowmik, Zachary Fox, Amit K Naskar, and John Gounley.  
 912 Transferring a molecular foundation model for polymer property predictions. *Journal of Chemical*  
 913 *Information and Modeling*, 63(24):7689–7698, 2023b. 29

918 Xujun Zhang, Odin Zhang, Chao Shen, Wanglin Qu, Shicheng Chen, Hanqun Cao, Yu Kang, Zhe  
 919 Wang, Ercheng Wang, Jintu Zhang, et al. Efficient and accurate large library ligand docking with  
 920 karmadock. *Nature Computational Science*, 3(9):789–804, 2023c. 27  
 921

922 Yangtian Zhang, Huiyu Cai, Chence Shi, Bozitao Zhong, and Jian Tang. E3bind: An end-to-end  
 923 equivariant network for protein-ligand docking. *arXiv preprint arXiv:2210.06069*, 2022a. 27  
 924

925 Zuobai Zhang, Minghao Xu, Arian Jamasb, Vijil Vijil, Aurelie Lozano, Payel Das, and Jian Tang.  
 926 Protein representation learning by geometric structure pretraining. In *International Conference  
 927 on Machine Learning*, 2022b. 26

928 Zuobai Zhang, Minghao Xu, Aurelie Lozano, Vijil Chenthamarakshan, Payel Das, and Jian Tang.  
 929 Pre-training protein encoder via siamese sequence-structure diffusion trajectory prediction. In  
 930 *Annual Conference on Neural Information Processing Systems*, 2023d. 26  
 931

932 Jiale Zhao, Wanru Zhuang, Jia Song, Yaqi Li, and Shuqi Lu. Pre-training protein bi-level repres-  
 933 entation through span mask strategy on 3d protein chains. *arXiv preprint arXiv:2402.01481*, 2024.  
 934 7, 26

935 Zibo Zhao, Wen Liu, Xin Chen, Xianfang Zeng, Rui Wang, Pei Cheng, Bin Fu, Tao Chen, Gang Yu,  
 936 and Shenghua Gao. Michelangelo: Conditional 3d shape generation based on shape-image-text  
 937 aligned latent representation. *Advances in neural information processing systems*, 36:73969–  
 938 73982, 2023. 9

939 Chao Zhou, Yanan Zhang, Jiaxin Chen, and Di Huang. Octr: Octree-based transformer for 3d  
 940 object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern  
 941 recognition*, pages 5166–5175, 2023a. 8

942 Gengmo Zhou, Zhifeng Gao, Qiankun Ding, Hang Zheng, Hongteng Xu, Zhewei Wei, Linfeng  
 943 Zhang, and Guolin Ke. Uni-mol: A universal 3d molecular representation learning framework.  
 944 In *The Eleventh International Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=6K2RM6wVqKu>. 8, 9, 28, 29  
 945

946 Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel  
 947 diffusion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages  
 948 5826–5835, 2021. 7, 25  
 949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972  
973  
974  
975  
976  
977

# Appendix

978  
979

## Table of Contents

980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990

<b>A The Merits of using Octrees for 3D Generation</b>	<b>20</b>
<b>B Implementation Details</b>	<b>20</b>
B.1 Fine-Grained Atom Tokenization for Microscopic Structures . . . . .	20
B.2 Vector Quantized Tokenization for Macroscopic Structures . . . . .	21
B.3 Model Architecture . . . . .	21
B.4 Input Embedding and Positional Encoding . . . . .	22
B.5 Generation Heads . . . . .	22
B.6 Efficiency Optimizations . . . . .	22
<b>C Experiment Settings</b>	<b>23</b>
C.1 3D Small Molecule Generation . . . . .	23
C.2 Crystal Generation . . . . .	23
C.3 Macroscopic 3D Object Generation . . . . .	25
C.4 Protein Pocket Prediction . . . . .	26
C.5 Molecular Docking . . . . .	27
C.6 Molecular Property Prediction via Pretraining . . . . .	28
C.7 Polymer Property Prediction via Pretraining . . . . .	29
<b>D More Experiments</b>	<b>29</b>
D.1 Mutual Benefits of Generation and Understanding Tasks . . . . .	29
D.2 Ablation Study . . . . .	30
D.3 Inference Speed . . . . .	31
<b>E Illustration of the Generated Examples</b>	<b>32</b>
<b>F LLM Usage Details</b>	<b>34</b>

1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

1026 

## A THE MERITS OF USING OCTREES FOR 3D GENERATION

1028 Octrees offer a principled way to turn sparse 3D geometry into short, informative token sequences—exactly what an autoregressive (AR) model needs to scale across domains and resolutions. Compared with uniform voxel grids and point/atom lists, an octree (i) adapts to sparsity, 1029 (ii) preserves precise spatial locality, and (iii) supplies a natural coarse-to-fine generation order that 1030 dramatically simplifies next-token prediction.

- 1034 • **Token efficiency at scale.** Let the finest grid resolution be  $2^L$  per axis, and let  $N$  denote the 1035 number of non-empty leaf cells at level  $L - 1$ . A uniform  $2^L \times 2^L \times 2^L$  grid yields  $O(8^L)$  tokens 1036 regardless of sparsity. In contrast, an octree emits at most  $8N$  tokens per level, totaling  $\leq 8NL$  1037 across  $L$  levels (see Section 2.1). With our 2-level subtree compression (2LSC), a parent and its 1038 8 children are encoded as *one* 8-bit token, reducing the count by  $\approx 8\times$  to  $\leq N(L - 1)$ . For thin 1039 structures (e.g., molecular surfaces or macroscopic shells) where  $N$  scales roughly with surface 1040 area, the token complexity approaches  $O(M^2 \log M)$  rather than  $O(M^3)$  for a grid of side length 1041  $M = 2^L$ —a decisive advantage for large systems.
- 1042 • **Coarse-to-fine inductive bias.** The octree’s hierarchy (Figure 2) gives each token strong context: 1043 high-level occupancy constrains where fine detail can appear, and subsequent levels specialize 1044 only within occupied regions. This bias shrinks the search space early—occupancy first, details 1045 later—so the AR model solves a sequence of easier problems rather than one monolithic one.
- 1046 • **Stable, explicit positions for AR prediction.** Point- or atom-based sequences suffer from ordering 1047 ambiguities and unknown future positions. Octree nodes, however, have deterministic positions 1048 (cell centers) and levels, which we feed as positional signals. Combined with our masked 1049 next-token prediction (MNTP; Section 2.2), the model conditions on the *correct* target position 1050 before predicting content, avoiding the instability of “predict-where-then-what” pipelines.
- 1051 • **Precision where it matters.** Deepening the tree only where geometry exists allocates resolution 1052 adaptively. Our fine-grained “3D patch” tokens then capture sub-voxel attributes (e.g., atom type 1053 and in-cell coordinates for molecules, or VQ codes for macroscopic shapes), marrying lossless 1054 spatial scaffolding with rich local detail (Section 2.1).
- 1055 • **Small, well-posed classification tasks.** 2LSC transforms eight binary occupancy decisions into 1056 a single 256-way classification, improving statistical efficiency and reducing sequence length. 1057 Downstream heads predict small discrete/continuous targets (e.g., token type and in-cell offsets) 1058 conditioned on strong spatial priors, which is well suited to AR transformers.
- 1059 • **Unified across scales and modalities.** Because the same octree scaffolding applies to Å-scale 1060 atoms and meter-scale objects, Uni-3DAR uses one tokenizer and one AR model for generation 1061 and understanding across molecules, crystals, proteins, and macroscopic shapes (Section 2.3). 1062 This uniformity simplifies conditioning (e.g., sequences, PXRD, text) and multi-frame tasks 1063 without custom architectures.

1064 In sum, the octree representation yields shorter sequences, clearer positional signals, and a natural 1065 generation curriculum. Together with 2LSC and MNTP, it makes AR modeling practical and 1066 accurate for cross-scale 3D generation and understanding.

1067 

## B IMPLEMENTATION DETAILS

1069 This section outlines the technical details of our approach, covering the tokenization schemes for 1070 different scales, the model architecture, and various optimizations.

1073 

### B.1 FINE-GRAINED ATOM TOKENIZATION FOR MICROSCOPIC STRUCTURES

1074 For microscopic 3D structures like molecules, we employ a fine-grained tokenization strategy where 1075 each token represents a single atom. This is achieved by recursively partitioning the 3D space using 1076 an octree until the final-level 3D patches are small enough to contain at most one atom. In our 1077 experiments, we set this final cell size,  $c_{L-1}$ , to  $0.24\text{\AA}$ .

1078 Each atom is thus represented by a token  $(t_i, e_i)$ , where  $t_i$  is the atom type (e.g., Carbon, Oxygen) 1079 and  $e_i = (e_i^0, e_i^1, e_i^2)$  specifies the atom’s coordinates within its cell (we don’t model the radius of a

1080 atom). To handle continuous positions, we discretize the coordinates with a resolution  $c_r = 0.01\text{\AA}$ ,  
 1081 mapping them to integers in the range  $\{0, \dots, N_p - 1\}$ , where  $N_p = c_{L-1}/c_r$ . In contrast, tokens  
 1082 representing non-terminal octree cells, which do not have a specific in-cell position, are assigned a  
 1083 default coordinate  $e_i = (N_p/2, N_p/2, N_p/2)$ . For data augmentation, we apply a random rotation  
 1084 to the structure before tokenization.

1085 This octree-based approach is highly efficient. For example, when applied to the QM9 dataset  
 1086 (Ramakrishnan et al., 2014a) using  $L = 6$  levels, a typical structure with an average of 18  
 1087 atoms is converted into approximately 160 tokens. This is a dramatic reduction compared to the  
 1088  $(2^6)^3 = 262,144$  tokens that would be required by a uniform grid of the same resolution. For other  
 1089 microscopic tasks, we keep the same size  $c_{L-1} = 0.24\text{\AA}$  for 3D patch, while the number of levels  
 1090  $L$  is set according to the data type. For example, we use  $L = 10$  for large proteins.  
 1091

## 1092 B.2 VECTOR QUANTIZED TOKENIZATION FOR MACROSCOPIC STRUCTURES

1093 For large, macroscopic 3D structures, we adopt a voxel-based representation and employ a Vector-  
 1094 Quantized Variational Autoencoder (VQ-VAE) for tokenization. This approach is analogous to  
 1095 methods used for 2D image tokenization, where an image is converted into a sequence of discrete  
 1096 tokens.  
 1097

1098 The core idea is to divide a high-resolution boolean voxel grid (e.g.,  $512 \times 512 \times 512$ ) into non-  
 1099 overlapping 3D patches and learn a discrete, compressed representation for each one. From the  
 1100 input grid resolution of  $512^3$  and a target latent grid of  $16^3$  tokens, each token ultimately represents  
 1101 a  $32 \times 32 \times 32$  patch of the original structure. To maintain a unified token format  $(t_i, e_i)$  with our  
 1102 other representations, the discrete code index from the VQ-VAE serves as the token type  $t_i$ , while  
 1103 its in-cell coordinate  $e_i$  is set to a default value.  
 1104

Our VQ-VAE tokenization pipeline involves the following steps:

1. **Lossless Voxel-to-Channel Packing:** We first perform a lossless pre-processing step to make the data more amenable to standard 3D convolutional networks. Each non-overlapping  $4 \times 4 \times 4$  block of the boolean input grid, containing 64 bits of information, is bit-packed into 8 bytes. This transforms the input data from a sparse, single-channel boolean tensor of shape  $1 \times 512^3$  into a dense, multi-channel tensor of shape  $8 \times 128^3$ , where each value is an integer in  $\{0, \dots, 255\}$ . This can be expressed as a mapping:  $\mathbb{B}^{1 \times 512 \times 512 \times 512} \rightarrow \mathbb{U}_8^{8 \times 128 \times 128 \times 128}$ .
2. **VQ-VAE Encoding:** A 3D VQ-VAE is trained on this  $8 \times 128 \times 128 \times 128$  multi-channel representation. The VQ-VAE's encoder network processes this volume using a downsampling factor of 8, mapping each  $8 \times 8 \times 8$  spatial patch of the multi-channel data to a single latent vector. This results in a final latent grid of  $16 \times 16 \times 16$  vectors. Each vector is then quantized by finding the nearest entry in a learned codebook.
3. **Token Representation:** The output of this process is a grid of integer indices,  $Z \in \{0, \dots, N_c\}^{16 \times 16 \times 16}$ , where each index corresponds to a vector in the codebook. Based on the provided code, we use a codebook with  $N_c = 512$  learnable "content" codes, where each code is a vector of dimensionality  $D_c = 4$ .

1121 To efficiently handle the inherent sparsity of most macroscopic structures, we introduce a special  
 1122 "blank" token. A  $32 \times 32 \times 32$  patch in the original voxel grid is considered blank if and only if all  
 1123 voxels within it are zero. During encoding, these blank patches are mapped to a reserved index (e.g.,  
 1124 index 0). The remaining  $N_c$  indices are used for non-empty patches. This allows subsequent generative  
 1125 models to ignore the blank tokens, focusing computational resources exclusively on regions  
 1126 containing geometry. We implement this VQ-VAE using the `vector-quantize-pytorch` library,  
 1127 configuring it with techniques like cosine similarity, k-means initialization, and diversity losses to  
 1128 ensure robust codebook utilization.  
 1129

## 1130 B.3 MODEL ARCHITECTURE

1131 We use a standard decoder-only Transformer architecture (Vaswani, 2017), based on the GPT-2  
 1132 model size. The model consists of 12 layers, an embedding dimension of 768, and 12 attention  
 1133 heads with a head dimension of 64, totaling approximately 90M parameters. Each layer contains a

---

1134           **Algorithm 1** A Simple Autoregressive Head for Sequential Target Prediction

---

1135           **Require:** Input tensor  $x$ , number of targets  $n$ , prediction heads for each target  $\{pred\_heads\}$ , embedding  
 1136            layers for each prediction  $\{emb\_layers\}$

1137            1:  $y \leftarrow x$   
 1138            2: Initialize  $preds \leftarrow \{\}$   
 1139            3: **for**  $i \leftarrow 1$  **to**  $n$  **do**  
 1140            4:     $p \leftarrow pred\_heads[i](y)$   
 1141            5:    Append  $p$  to  $preds$   
 1142            6:     $y \leftarrow y + emb\_layers[i](p)$  {Teacher-forcing during training}  
 1143            7: **end for**  
 1144            8: **return**  $preds$

---

1145  
 1146            unidirectional self-attention module and a SwiGLU (Shazeer, 2020) feed-forward network. For nor-  
 1147            malization, we employ a pre-norm design (Xiong et al., 2020) with RMSNorm (Zhang and Sennrich,  
 1148            2019).

#### 1150           B.4 INPUT EMBEDDING AND POSITIONAL ENCODING

1151  
 1152            The input representation for the  $i$ -th token combines several pieces of information: its type  $t_i$ , in-  
 1153            cell coordinates  $e_i$ , octree level  $l_i$ , frame index  $f_i$  (for multi-frame sequences), and its absolute 3D  
 1154            coordinate  $c_i$ . For octree and masked tokens,  $c_i$  is the center of the corresponding cell. For atom  
 1155            tokens, we use the precise atom coordinate for  $c_i$  to provide a more accurate positional signal.

1156            These discrete attributes  $(t_i, e_i, l_i, f_i)$  are converted into high-dimensional vectors via separate em-  
 1157            bedding layers, and their embeddings are summed to form the final input to the model. Notably,  
 1158            our method does not use any 2D graphical information, such as chemical bonds, making it broadly  
 1159            applicable to diverse 3D data. For encoding pairwise positional information, we apply 3D Rotary  
 1160            Position Embedding (RoPE-3D) (Su et al., 2024) to the absolute coordinates  $c_i$ .

#### 1162           B.5 GENERATION HEADS

1163  
 1164            The model’s generative task is to predict the content of masked tokens. For an **octree token**, only the  
 1165            type  $t_i$  needs to be predicted (since  $e_i$  is fixed), which is handled by a simple classification head. For  
 1166            an **atom token**, both the type  $t_i$  and the in-cell coordinates  $e_i$  must be predicted. After predicting  
 1167             $t_i$ , we predict  $e_i$  using one of two methods:

- 1168            • **Autoregressive Prediction:** The coordinates  $(e_i^0, e_i^1, e_i^2)$  are predicted sequentially, as detailed in  
 1169            Alg. 1.
- 1170            • **Diffusion Prediction:** We adapt the token-level diffusion module from MAR (Li et al., 2024b) to  
 1171            generate the continuous coordinates  $e_i$ .

1172            Our experiments showed that both methods yield similar performance (see Sec. D.2). We therefore  
 1173            use the more computationally efficient autoregressive approach as our default. During inference,  
 1174            we employ a sampling strategy to balance quality and diversity: we first sample from the model  
 1175            using a slightly elevated temperature and then select the top- $r$  candidates based on their cumulative  
 1176            autoregressive probabilities. This method has proven more effective than standard low-temperature  
 1177            sampling.

#### 1179           B.6 EFFICIENCY OPTIMIZATIONS

1180  
 1181            We implement several optimizations to ensure efficient training and inference. **During training**, we  
 1182            use FlashAttention (Dao et al., 2022) with bfloat16 to accelerate computation and reduce memory  
 1183            usage. We also employ sequence packing, where tokens from multiple samples are concatenated into  
 1184            a single sequence. This technique eliminates the overhead of padding and is particularly effective for  
 1185            handling systems of varying sizes, such as proteins. **During inference**, we use a KV-cache to speed  
 1186            up token generation. To further improve throughput for masked prediction, we generate tokens in  
 1187            pairs instead of one by one. This is possible because the inputs for masked tokens are known in  
 1188            advance, allowing us to pack adjacent prediction steps to better utilize the GPU.

1188 Table 5: Our experiments cover a broad spectrum of real-world tasks, each of which can be seamlessly adapted  
 1189 by the unified framework of Uni-3DAR.

Section	Data Type	Single-Frame Gen.	Multi-Frame Gen.	Token Und.	Structure Und.
Sec. 3.1	Molecule	✓			
Sec. 3.2	Crystal + PXRD		✓		
Sec. 3.3	Macroscopic 3D Object	✓			
Sec. 3.4	Protein	✓		✓	
Sec. 3.5	Protein + Molecule		✓		
Sec. 3.6	Molecule / Polymer	✓			✓

## C EXPERIMENT SETTINGS

### C.1 3D SMALL MOLECULE GENERATION

Generating small organic molecules with accurate 3D conformations is a classical, benchmark-rich task in molecular modeling, yet the inherent flexibility due to rotatable bonds and diverse conformations poses significant challenges. Evaluating Uni-3DAR on this task directly tests its capability to generate realistic 3D molecular structures through a straightforward application of its single-frame generation methodology.

**Dataset and Metric** Consistent with previous studies (Hoogeboom et al., 2022), we use the QM9 (Ramakrishnan et al., 2014b) and GEOM-DRUG (Axelrod and Gomez-Bombarelli, 2022) datasets for unconditional 3D molecular generation. QM9, a widely-used molecular machine learning benchmark, contains 130K small molecules with high-quality 3D conformations (up to 9 heavy atoms and 29 total atoms including hydrogens), split into training (100K), validation (18K), and test sets (13K). GEOM-DRUG, in contrast, features larger organic compounds containing up to 181 atoms (averaging 44.2 atoms across 5 types), covering approximately 37 million conformations for around 450K unique molecules. Following established protocols (Hoogeboom et al., 2022), we select the 30 lowest-energy conformations per molecule for training.

Model performance is evaluated based on chemical feasibility. Bond types (single, double, triple, or none) are inferred from molecular geometries using pairwise atomic distances and atom types. Metrics include Atom Stability (the fraction of atoms exhibiting correct valency), Molecule Stability (the percentage of molecules where all atoms are stable), validity (percentage of chemically valid molecules verified by RDKit), and uniqueness (percentage of unique compounds among generated molecules). Metrics are computed consistently using the evaluation code from previous studies (Hoogeboom et al., 2022).

**Baselines and Implementation** We benchmark Uni-3DAR against established models, including G-SchNet (Gebauer et al., 2022), ENF (Garcia Satorras et al., 2021), EDM (Hoogeboom et al., 2022) and its variants GDM (Hoogeboom et al., 2022), EDM-Bridge (Wu et al., 2022), GeoLDM (Xu et al., 2023b), and UniGEM (Feng et al., 2024), which uses additional molecular properties to enhance generation performance.

Uni-3DAR employs a single-frame generation approach with a batch size of 64 for QM9 and 128 for GEOM-DRUG. The model is trained for 500K steps (approximately 320 epochs for QM9 and 12 epochs for GEOM-DRUG). We apply a peak learning rate of 3e-4, incorporating a 6% linear warmup phase followed by cosine decay. Training duration is approximately 6.9 hours on 4 NVIDIA 4090 GPUs for QM9 and around 11.7 hours on 8 NVIDIA 4090 GPUs for GEOM-DRUG.

### C.2 CRYSTAL GENERATION

**Tasks** Unlike organic molecules, crystal structures are typically rigid with stable conformations. However, crystals introduce unique challenges due to their inherent symmetry and periodic arrangement in 3D space. A crystal is conventionally represented by its lattice (a parallelepiped unit cell) along with atomic details, including atom types and their coordinates within the lattice.

1242 In Uni-3DAR, crystal structure generation is approached as a two-frame generative process: first  
 1243 generating the eight vertices defining the lattice, followed by generating the atomic configurations  
 1244 inside the generated lattice. Notably, unlike previous methods employing fractional coordinates, we  
 1245 consistently use physical coordinates to maintain uniformity across various molecular data types.  
 1246

1247 Based on this generation approach, we define and address three distinct tasks:

- 1248 1. *De Novo Crystal Generation*: Learning the distribution of crystal structures from data to generate  
 1249 novel samples unconditionally.
- 1250 2. *Crystal Structure Prediction (CSP)*: Predicting crystal structures from given chemical composi-  
 1251 tions (atom types and counts). During inference, the chemical composition is provided as condi-  
 1252 tion, enabling the model to generate the corresponding crystal structure.
- 1253 3. *PXRD-guided Crystal Structure Prediction*: Establishing a cross-modal mapping from powder  
 1254 X-ray diffraction (PXRD) signals and chemical compositions to reconstruct crystal structures that  
 1255 accurately match observed PXRD patterns. This task has significant practical implications, as  
 1256 PXRD analysis is widely used in crystal structure determination and validation of novel materials  
 1257 in real-world scenarios.

1259  
 1260 **Dataset and Metric** We employ established datasets consistent with prior studies (Xie et al., 2021;  
 1261 Jiao et al., 2023; Miller et al., 2024) for both training and evaluation purposes. Specifically, we  
 1262 employ the Carbon-24 dataset (Pickard, 2020), containing 10,153 carbon-based structures with cells  
 1263 composed of 6 to 24 atoms. The MP-20 dataset (Jain et al., 2013), derived from the Materials  
 1264 Project (Jain et al., 2013), includes 45,231 stable inorganic materials representing a wide range of  
 1265 experimentally validated compounds, each containing up to 20 atoms per cell. Additionally, we use  
 1266 the more challenging MPTS-52 dataset, an extended version of MP-20, comprising 40,476 structures  
 1267 with up to 52 atoms per cell, organized by the earliest publication year. We follow the same data  
 1268 split strategy as outlined in previous work (Jiao et al., 2023).

1269 To evaluate de novo crystal generation performance, we adopt the standard evaluation framework  
 1270 proposed by Xie et al. (2021), which includes three key metrics: validity, coverage, and property  
 1271 statistics. Validity quantifies the proportion of generated structures that satisfy established physical  
 1272 plausibility criteria. Coverage measures the ability of generated structures to capture the diversity  
 1273 present in the test set. Property statistics compare essential attributes such as density, formation  
 1274 energy, and elemental composition between generated and ground-truth distributions.

1275 For assessing performance in CSP and PXRD-guided CSP tasks, we align our evaluation method-  
 1276 ology with prior research (Miller et al., 2024). We compute the top-1 match rate alongside  
 1277 the corresponding average root-mean-square error (RMSE) for matched structures. We employ  
 1278 StructureMatcher(Ong et al., 2013), using thresholds set to `stol`=0.5, `angle_tol`=10, and  
 1279 `ltol`=0.3, consistent with the methodology of previous studies (Miller et al., 2024).

1280  
 1281 **Baseline Models and Implementation** We benchmark Uni-3DAR against established methods,  
 1282 including FTCP (Ren et al., 2021), G-SchNet (Gebauer et al., 2019), P-G-SchNet (Gebauer et al.,  
 1283 2019), CDVAE (Xie et al., 2021), DiffCSP (Jiao et al., 2023), and FlowMM (Miller et al., 2024).  
 1284 Additionally, we evaluate Uni-3DAR against the recent UniGenX (Zhang et al., 2025) for the CSP  
 1285 task. For PXRD-guided CSP, we compare Uni-3DAR with PXRDGEN (Li et al., 2024a), a model  
 1286 tailored for this task.

1287 In Uni-3DAR, we use a 12-layer model with a 768-dimensional embedding for de novo crystal  
 1288 generation, while a larger 24-layer model with a 1024-dimensional embedding is employed for CSP  
 1289 and PXRD-guided CSP tasks. All models are trained for 400k steps with a batch size of 64 and a  
 1290 peak learning rate of 3e-4. For chemical composition conditioning, we prepend a token derived from  
 1291 a multi-hot atom-type vector. PXRD data, spanning angles from 0° to 120°, is converted into a 1200-  
 1292 dimensional vector with a 0.1° resolution, evenly divided into four segments, each represented by a  
 1293 conditional token. As a result, PXRD-guided CSP utilizes a total of five conditional tokens (one for  
 1294 composition and four for PXRD signals). The autoregressive nature of Uni-3DAR enables seamless  
 1295 integration of these conditional tokens, eliminating the need for additional encoders required by  
 1296 previous methods (Li et al., 2024a; Lai et al., 2025).

1296 **Results of De Novo Crystal Generation** The performance of Uni-3DAR on the Carbon-24 and  
 1297 MP-20 datasets is presented in Table 2. On Carbon-24, Uni-3DAR outperforms existing models, par-  
 1298 ticularly excelling in coverage, demonstrating its ability to generate diverse and realistic structures.  
 1299 On MP-20, Uni-3DAR significantly enhances component validity compared to previous approaches  
 1300 while maintaining competitive performance on other metrics. These results underscore Uni-3DAR’s  
 1301 strength in producing chemically valid crystal structures that closely align with key physical and  
 1302 chemical properties.

1303 **Results of Crystal Structure Prediction (CSP)** We evaluate Uni-3DAR’s performance on CSP  
 1304 across all datasets, as summarized in Table 3. Uni-3DAR consistently outperforms baseline methods  
 1305 by significant margins. Specifically, on Carbon-24, it improves the match rate by 4.14% over the  
 1306 previous best method, demonstrating superior accuracy in reconstructing crystal structures. On MP-  
 1307 20, Uni-3DAR achieves a substantial improvement in RMSE, reducing it from 0.0566 to 0.0317,  
 1308 a relative improvement of 178% over the second-best model. Furthermore, on MPTS-52, Uni-  
 1309 3DAR achieves an impressively low RMSE of 0.0684, representing a 184% relative improvement,  
 1310 despite the increased structural complexity. This result highlights its exceptional precision in atomic  
 1311 placement. Overall, these findings demonstrate Uni-3DAR’s strong generalization capability across  
 1312 datasets of varying difficulty levels.

1313 **Results of PXRD-Guided CSP** Table 3 demonstrates Uni-3DAR’s performance in PXRD-guided  
 1314 CSP on the MP-20 dataset, benchmarked against PXRDGEN (Li et al., 2024a). Uni-3DAR substan-  
 1315 tially outperforms PXRDGEN, elevating the match rate from 68.68% to 75.08% while drastically  
 1316 reducing the RMSE from 0.0707 to 0.0276—a 256% relative improvement. This significant RMSE  
 1317 reduction underscores Uni-3DAR’s exceptional ability to generate crystal structures that precisely  
 1318 correspond to experimental PXRD patterns. Collectively, these results underscore the superior ca-  
 1319 pability of Uni-3DAR in harnessing diffraction constraints to reliably predict crystal structures.

### 1321 C.3 MACROSCOPIC 3D OBJECT GENERATION

1322 To demonstrate versatility beyond the microscopic realm of molecules and crystals, we further eval-  
 1323 uate Uni-3DAR on unconditional macroscopic 3D object generation, a core task in 3D computer  
 1324 vision (3DCV). The goal is to synthesize realistic and diverse 3D shapes of everyday objects di-  
 1325 rectly from the learned data distribution.

1326 **Dataset and Evaluation Protocol** Following common practice, we adopt three categories from  
 1327 ShapeNet (Chang et al., 2015)—airplane, chair, and car. Each object is represented as a point  
 1328 cloud with 2,048 points uniformly sampled from the surface. In line with recent recommendations,  
 1329 we evaluate using **1-nearest-neighbor accuracy (1-NNA; lower is better)** computed with both  
 1330 Chamfer Distance (CD) and Earth Mover’s Distance (EMD) (Yang et al., 2019; Vahdat et al., 2022).  
 1331 Concretely, given a generated set  $S_g$  and a reference set  $S_r$ , 1-NNA is the leave-one-out accuracy  
 1332 of a 1-NN classifier on  $S_g \cup S_r$ ; if  $S_g$  matches  $S_r$  well, the classification accuracy approaches  
 1333 50%. Compared with legacy metrics such as coverage (COV) and minimum matching distance  
 1334 (MMD), 1-NNA more directly captures distributional similarity while jointly reflecting both quality  
 1335 and diversity, and avoids several known failure modes of COV/MMD. We therefore report 1-NNA  
 1336 (with CD/EMD) as our primary metric throughout this section and in the main paper.

1337 **Baselines and Implementation Details** We benchmark Uni-3DAR against established point-  
 1338 cloud generative models, including r-GAN and l-GAN (Achlioptas et al., 2018), PointFlow (Yang  
 1339 et al., 2019), SoftFlow (Kim et al., 2020), SetVAE (Kim et al., 2021), DPF-Net (Klokov et al.,  
 1340 2020), diffusion-based methods DPM (Luo and Hu, 2021) and PVD (Zhou et al., 2021), and the  
 1341 recent LION (Vahdat et al., 2022). To ensure clear and reproducible comparisons, we follow the  
 1342 PointFlow data protocol and training/test splits for the three categories.

1343 For Uni-3DAR, the input 3D object is voxelized at  $512 \times 512 \times 512$  resolution. We define fine-  
 1344 grained structural tokens as non-overlapping  $16 \times 16 \times 16$  voxel patches and quantize each patch  
 1345 with a VQVAE codebook. (In our main text we summarize this as “each patch is quantized using  
 1346 VQVAE”; here we provide the fuller setup for completeness.) Unless otherwise specified, generation  
 1347 uses our single-frame sampling procedure, analogous to the molecular setting. For each ShapeNet  
 1348 category, the VQVAE is trained for 200 epochs; Uni-3DAR is then trained for 10,000 steps with

1350  
1351 Table 6: Results for atom-level binding site prediction measured by IoU (%). Baseline results are taken from  
1352 [Zhao et al. \(2024\)](#). For a fair comparison with other methods, we report Vabs-Net’s result using only  $\alpha$ -carbon  
1353 atoms.

Method	pretrained	B277↑	DT198↑	ASTEX85↑	CHEN251↑	COACH420↑
FPocket ( <a href="#">Le Guilloux et al., 2009</a> )	×	31.5	23.2	34.1	25.4	30.0
SiteHound ( <a href="#">Hernandez et al., 2009</a> )	×	36.4	23.1	38.9	29.4	34.9
MetaPocket2 ( <a href="#">Macari et al., 2019</a> )	×	37.3	25.8	37.5	32.8	37.7
DeepSite ( <a href="#">Jiménez et al., 2017</a> )	×	34.0	29.1	37.4	27.4	33.9
P2Rank ( <a href="#">Krivák and Hoksza, 2018</a> )	×	<b>49.8</b>	<b>38.6</b>	<b>47.4</b>	<b>56.5</b>	45.3
ESM2_150M ( <a href="#">Lin et al., 2023b</a> )	✓	19.6	16.6	20.5	18.9	22.0
GearNet ( <a href="#">Zhang et al., 2022b</a> )	✓	39.9	35.8	41.0	36.4	41.3
Siamdiff ( <a href="#">Zhang et al., 2023d</a> )	✓	37.7	31.0	40.7	35.3	40.3
Vabs-Net ( <a href="#">Zhao et al., 2024</a> )	✓	-	-	-	-	<b>56.3</b>
Uni-3DAR	✓	<b>53.4</b>	<b>46.7</b>	<b>51.4</b>	<b>47.9</b>	<b>56.2</b>

1363  
1364 batch size 64. On a single NVIDIA RTX 4090, training per category requires approximately 10  
1365 hours for the VQVAE and 2 hours for Uni-3DAR.

1366  
1367  
1368 **Results** Table 4 (main paper) summarizes unconditional generation under the 1-NNA protocol.  
1369 Uni-3DAR achieves the lowest (best) 1-NNA in all six category–metric pairs (Airplane/Chair/Car  
1370  $\times$  CD/EMD), outperforming strong diffusion and flow-based baselines. In particular, Uni-3DAR  
1371 consistently improves over LION—the strongest baseline in our comparison—by small but system-  
1372 atic margins: *Airplane* (CD: 67.35 vs. 67.41; EMD: 61.09 vs. 61.23), *Chair* (CD: 53.11 vs. 53.70;  
1373 EMD: 50.98 vs. 52.34), and *Car* (CD: 53.35 vs. 53.41; EMD: 50.89 vs. 51.14). Taken together, these  
1374 results indicate that Uni-3DAR produces point-cloud distributions that are both high-quality and di-  
1375 verse, closely matching the real data according to a metric expressly designed to assess distributional  
1376 similarity.

#### 1377 1378 C.4 PROTEIN POCKET PREDICTION

1379  
1380 Proteins are a crucial class of biological structures, and accurate prediction of binding pockets is  
1381 essential for de novo drug design and applications such as molecular docking. Traditionally, pocket  
1382 prediction is formulated as an atom-level or residue-level classification task. Each atom or residue  
1383 is assigned a binary label indicating whether it belongs to a binding pocket. We adopt this classical  
1384 formulation to evaluate Uni-3DAR’s token-level understanding capabilities.

1385  
1386  
1387 **Dataset and Metric** We follow previous studies ([Zhao et al., 2024](#)) and employ a binding site  
1388 dataset constructed from the CASF-2016 core set ([Su et al., 2018](#)), PDBBind v2020 refined set ([pdb](#),  
1389 [2025](#)), and MOAD ([Hu et al., 2005](#)). The dataset consists of 23k training samples, 5k validation  
1390 samples, and five test sets of roughly 1k samples each. Model performance is assessed using the  
1391 Intersection-over-Union (IoU) metric, consistent with previous evaluations ([Zhao et al., 2024](#)).

1392  
1393  
1394 **Baselines and Implementation** We benchmark Uni-3DAR against established methods. Our  
1395 comparisons include non-pretrained approaches (e.g., FPocket ([Le Guilloux et al., 2009](#)), SiteHound  
1396 ([Hernandez et al., 2009](#)), etc.) and pretrained models (e.g., ESM2\_150M ([Lin et al., 2023b](#)), GearNet  
1397 ([Zhang et al., 2022b](#)), Siamdiff ([Zhang et al., 2023d](#)), and Vabs-Net ([Zhao et al., 2024](#))). In line with  
1398 prior works ([Zhao et al., 2024](#)), we pretrain Uni-3DAR on approximately 1.3 million protein struc-  
1399 tures before fine-tuning it on the binding site dataset. Unlike Vabs-Net, which employs full-atom  
1400 representations, our experiments are restricted to  $\alpha$ -carbon atoms to facilitate direct comparisons.

1401  
1402 Pretraining is conducted using a single-frame generation approach for 300k steps with a batch size  
1403 of 64. We use a peak learning rate of 3e-4 with a 10% linear warmup followed by cosine decay,  
which requires approximately 19 hours on 16 NVIDIA A100 GPUs. Fine-tuning adopts an atom-  
1404 level classification strategy, conducted for 100 epochs with a batch size of 32, a peak learning rate  
1405 of 1e-4, requiring roughly 7 hours on 8 NVIDIA A100 GPUs.

Table 7: Comparison of docking performance on the Top1- and Top5-RMSD metrics. The first group of five baselines comprises classical docking software, while the second group of eight baselines consists of deep learning-based methods. The results are reproduced directly from Cao et al. (2024). The best outcomes are shown in **bold**, and the second-best are underlined.

		Top1-RMSD			Top5-RMSD		
		%<1Å ↑	%<2Å ↑	Med(Å) ↓	%<1Å ↑	%<2Å ↑	Med(Å) ↓
Uni-Dock (Yu et al., 2022)		32.51±0.39	50.69±0.59	1.89±0.04	47.11±0.22	67.03±0.94	1.10±0.02
Glide SP (Friesner et al., 2004)		17.36±0.00	44.63±0.00	2.27±0.00	31.13±0.00	60.06±0.00	1.54±0.00
GNINA (Ragoza et al., 2017)		21.12±0.26	43.62±1.06	2.45±0.07	28.47±0.57	58.13±0.81	1.65±0.02
SMINA (Koes et al., 2013)		18.73±0.00	31.68±0.00	3.99±0.00	28.47±0.56	48.48±0.00	2.07±0.00
Vina (Eberhardt et al., 2021)		18.32±0.02	36.64±0.05	3.42±0.01	24.79±0.00	50.96±0.00	1.87±0.01
EquiBind (Stärk et al., 2022b)		/	5.5±1.2	6.2±0.3	/	/	/
TANKBind (Lu et al., 2022)		2.66±0.26	18.18±0.60	4.2±0.05	4.13±0.0	20.39±0.45	3.5±0.04
E3Bind (Zhang et al., 2022a)		/	25.6	7.2	/	/	/
KarmaDock (Zhang et al., 2023c)		/	56.2	/	/	/	/
DiffDock(Pocket) (Corso et al.)		/	51.8	2.0	/	60.7	1.9
DiffDock (Corso et al.)		15.15	36.09	3.35	21.76	43.52	2.46
DiffDock-L (Corso et al., 2024)		19.07±0.57	40.74±1.25	2.88±0.18	21.95±0.39	48.15±0.91	2.05±0.04
SurfDock (Cao et al., 2024)		40.96±0.34	<u>68.41</u> ±0.26	<u>1.18</u> ±0.00	<u>54.18</u> ±0.13	<b>75.11</b> ±0.13	<u>0.94</u> ±0.00
Uni-3DAR		<b>44.75</b> ±2.63	<b>69.06</b> ±0.75	<b>1.08</b> ±0.04	<b>56.35</b> ±1.99	<u>72.38</u> ±0.73	<b>0.76</b> ±0.02

## C.5 MOLECULAR DOCKING

Molecular docking predicts how a ligand binds to a target protein, playing a crucial role in drug discovery. In Uni-3DAR, this process is structured as a three-frame generation task. The first two frames represent the protein and the initial ligand, both provided as inputs during inference, while the third frame corresponds to the predicted docked conformation of the ligand.

**Dataset and Metric** Following Cao et al. (2024), we train and evaluate docking methods on the PDBbind2020 dataset. The training and validation set consists of 17,000 complexes from 2018 or earlier, while the test set includes 363 structures from 2019, ensuring no ligand overlap with the training data. Given a protein-binding pocket and a randomly generated ligand conformation from RDKit, the goal is to generate a user-specified number of poses (set to 40, as in Cao et al. (2024)). Docking methods typically incorporate a confidence scoring mechanism to rank these poses. Performance is assessed using the percentage of predictions with RMSD < 1Å and RMSD < 2Å, as well as the median RMSD for the top-ranked pose and the best pose among the top five ranked poses.

**Baselines and Implementation** We evaluate Uni-3DAR against 13 baselines, including five classical docking software tools and eight deep learning-based methods. Most existing deep learning approaches rely on complex featurizations, such as using protein language model embeddings (e.g., from ESM2 (Lin et al., 2022)). To simplify and unify the molecular tasks, we omit these complicated features in Uni-3DAR and instead use only atom types and coordinates. We also adopt a full-atom representation of the protein pocket to enhance expressive power. We frame docking as an autoregressive generation task by embedding both the pocket and the RDKit conformation as two frames, concatenating them into a single input sequence, and training the model to generate the docked molecule conformation as a new frame sequence. For further simplicity, we do not impose constraints such as matching the number and types of atoms in the output frame to those of the input molecule. Also, we do not train a separate scoring model for pose ranking. Instead, we use the cumulative probability derived from autoregressive generation to score each generated pose. We train Uni-3DAR for 300k steps (approximately 300 epochs) with a batch size of 16. The learning rate schedule follows the same configuration as the experiments detailed in section 3.1. The training is completed in approximately one day on 4 NVIDIA A100 GPUs.

**Results** Experimental results are summarized in Table 7. Uni-3DAR outperforms the state-of-the-art method, SurfDock, demonstrating similar percentages of poses with RMSD below 1Å and 2Å. Notably, Uni-3DAR excels in generating higher-quality poses, reflected by its lower median

1458 Table 8: Results on molecular property prediction performance. The best results are highlighted in **bold**, and  
 1459 the second-best results are underlined. Baseline results are taken from [Lu et al. \(2025\)](#).

Model	HOMO $\downarrow$	LUMO $\downarrow$	GAP $\downarrow$	E1-CC2 $\downarrow$	E2-CC2 $\downarrow$	f1-CC2 $\downarrow$	f2-CC2 $\downarrow$	Dipmom $\downarrow$	aIP $\downarrow$	D3_disp $\downarrow$
	(Hartree)	(Hartree)	(Hartree)	(eV)	(eV)	(Debye)	(eV)	(eV)	_corr (eV)	
GROVER (Rong et al., 2020)	0.0075	0.0086	0.0109	0.0101	0.0129	0.0219	0.0401	0.0752	0.1467	0.2516
	$\pm 2.0e-4$	$\pm 8.0e-4$	$\pm 1.4e-3$	$\pm 9.7e-4$	$\pm 4.6e-4$	$\pm 3.5e-4$	$\pm 1.2e-3$	$\pm 1.1e-3$	$\pm 1.5e-2$	$\pm 5.3e-2$
GEM (Fang et al., 2022)	0.0068	0.0080	0.0107	0.0090	0.0102	0.0170	0.0352	0.0289	0.0207	0.0077
	$\pm 7.0e-5$	$\pm 2.0e-5$	$\pm 1.9e-4$	$\pm 1.3e-4$	$\pm 2.3e-4$	$\pm 4.3e-4$	$\pm 5.4e-4$	$\pm 1.2e-3$	$\pm 2.6e-4$	$\pm 6.6e-4$
3D Infomax (Stärk et al., 2022a)	0.0065	0.0070	0.0095	0.0089	0.0091	0.0172	0.0364	0.0291	0.0526	0.2285
	$\pm 1.0e-5$	$\pm 1.0e-4$	$\pm 1.0e-4$	$\pm 2.0e-4$	$\pm 3.0e-4$	$\pm 4.0e-4$	$\pm 9.0e-4$	$\pm 1.7e-3$	$\pm 1.4e-4$	$\pm 7.5e-3$
Uni-Mol (Zhou et al., 2023b)	0.0052	0.0060	0.0081	0.0067	0.0080	0.0143	0.0309	<u>0.0106</u>	<u>0.0095</u>	<b>0.0047</b>
	$\pm 2.0e-5$	$\pm 6.0e-5$	$\pm 4.0e-5$	$\pm 4.0e-5$	$\pm 4.0e-5$	$\pm 2.0e-4$	$\pm 9.4e-4$	$\pm 3.1e-4$	$\pm 6.4e-4$	$\pm 5.6e-4$
Mol-AE (Yang et al., 2024)	0.0050	0.0057	0.0080	0.0070	0.0080	<u>0.0140</u>	0.0307	0.0113	0.0103	0.0077
	$\pm 8.0e-5$	$\pm 4.7e-4$	$\pm 8.0e-5$	$\pm 6.0e-5$	$\pm 4.0e-5$	$\pm 4.0e-5$	$\pm 1.3e-3$	$\pm 4.7e-4$	$\pm 1.3e-4$	$\pm 1.3e-3$
SpaceFormer (Lu et al., 2025)	<b>0.0042</b>	<b>0.0040</b>	<b>0.0064</b>	<u>0.0058</u>	<u>0.0074</u>	0.0142	<u>0.0294</u>	<b>0.0083</b>	<b>0.0090</b>	0.0053
	$\pm 1.0e-5$	$\pm 2.0e-5$	$\pm 1.2e-4$	$\pm 8.0e-5$	$\pm 8.4e-5$	$\pm 3.7e-4$	$\pm 7.1e-4$	$\pm 5.0e-4$	$\pm 5.9e-4$	$\pm 1.2e-3$
Uni-3DAR	<u>0.0048</u>	<u>0.0044</u>	<u>0.0065</u>	<b>0.0056</b>	<b>0.0067</b>	<b>0.0134</b>	<b>0.0286</b>	0.0114	0.0127	<u>0.0052</u>
	$\pm 2.1e-5$	$\pm 3.2e-5$	$\pm 8.8e-5$	$\pm 2.2e-5$	$\pm 2.0e-5$	$\pm 7.0e-5$	$\pm 1.6e-4$	$\pm 6.9e-4$	$\pm 1.1e-4$	$\pm 3.2e-4$

1473  
 1474 RMSD values. However, Uni-3DAR exhibits slightly inferior performance in selecting Top-5 poses  
 1475 for challenging cases, as evidenced by a lower percentage of poses with RMSD below 2 $\text{\AA}$  in the  
 1476 Top5-RMSD evaluation (72.38% vs. 75.11% for SurfDock). This gap may arise because the scoring  
 1477 module in Uni-3DAR has not been explicitly trained and is only exposed to ground-truth conforma-  
 1478 tions during the training phase. Addressing this limitation by training a dedicated scoring module  
 1479 could potentially enhance its selection performance. Moreover, since Uni-3DAR avoids complex  
 1480 feature engineering, its docking accuracy might further benefit from multitask learning strategies,  
 1481 emphasizing the promise of a unified foundational model for molecular applications.

## 1482 C.6 MOLECULAR PROPERTY PREDICTION VIA PRETRAINING

1483 Molecular property prediction through pretraining has emerged as an effective strategy to address  
 1484 data scarcity challenges in areas like drug discovery and material design. As a classical task with  
 1485 established benchmarks, molecular property prediction directly assesses a model’s capacity to com-  
 1486 prehend 3D molecular structures. Applying Uni-3DAR’s structure-level understanding framework  
 1487 is thus straightforward.

1488  
 1489 **Dataset and Metric** We utilize the same pretraining dataset as employed by Uni-Mol (Zhou et al.,  
 1490 2023b) and SpaceFormer (Lu et al., 2025), comprising approximately 19 million molecules. For  
 1491 downstream evaluations, we follow the datasets and evaluation settings used by the state-of-the-  
 1492 art SpaceFormer (Lu et al., 2025). These include a 20K dataset predicting electronic properties  
 1493 (HOMO, LUMO, GAP), a 21K dataset targeting energy properties (E1-CC2, E2-CC2, f1-  
 1494 CC2, f2-CC2), and an 8K dataset predicting mechanical and electronic properties (Dipmom, aIP, and D3  
 1495 Dispersion Corrections). Data splits align exactly with SpaceFormer’s methodology (Lu et al., 2025).  
 1496 Performance across all tasks is measured using the Mean Absolute Error (MAE) metric.

1497  
 1498 **Baselines and Implementation** Our baselines encompass several prominent models, including  
 1499 Uni-Mol (Zhou et al., 2023b), Mol-AE (Yang et al., 2024), 3D Infomax (Stärk et al., 2022a),  
 1500 GROVER (Rong et al., 2020), GEM (Fang et al., 2022), and the most recent state-of-the-art method,  
 1501 SpaceFormer (Lu et al., 2025). For pretraining, we use the proposed masked next-token prediction  
 1502 as pretraining task, training the model for 500k steps with a batch size of 128. The peak learning rate  
 1503 is set to 3e-4, incorporating a 10% linear warmup followed by cosine decay, requiring approximately  
 1504 11.5 hours on 8 NVIDIA 4090 GPUs.

1505  
 1506 During fine-tuning, we adopt a structure-level understanding strategy, supplemented by a masked  
 1507 next-token prediction auxiliary generative loss. Training is conducted over a maximum of 200  
 1508 epochs. We systematically explore hyperparameter combinations, considering two batch sizes (32,  
 1509 64) and two learning rates (5e-4, 1e-4), resulting in four distinct setups. For each hyperparameter  
 1510 configuration, models are trained three times using different random seeds, and we report the mean  
 1511 performance along with standard deviation. The best-performing model based on validation loss is  
 selected for evaluation.

1512 Table 9: Polymer properties prediction performance. The best results are highlighted in **bold**, and the second-  
1513 best results are underlined.

Model	Egc ↓ (eV)	Egb ↓ (eV)	Eea ↓ (eV)	Ei ↓ (eV)	Xc ↓ %	Eps ↓ 1	Nc ↓ 1	Eat ↓ eV/atom
ChemBERTa (Chithrananda et al., 2020)	0.539 ± 0.049	0.664 ± 0.079	0.350 ± 0.036	0.485 ± 0.086	18.711 ± 1.396	0.603 ± 0.083	0.140 ± 0.010	0.219 ± 0.056
Uni-Mol (Zhou et al., 2023b)	0.489 ± 0.028	0.531 ± 0.055	0.332 ± 0.027	0.407 ± 0.080	17.414 ± 1.581	0.536 ± 0.053	0.095 ± 0.013	0.084 ± 0.034
SML (Zhang et al., 2023b)	0.489 ± 0.056	0.547 ± 0.110	0.313 ± 0.016	0.432 ± 0.060	18.981 ± 1.258	0.576 ± 0.020	0.102 ± 0.010	0.062 ± 0.014
PLM (Zhang et al., 2023b)	0.459 ± 0.036	0.528 ± 0.081	0.322 ± 0.037	0.444 ± 0.062	19.181 ± 1.308	0.576 ± 0.060	0.100 ± 0.010	<b>0.050</b> ± 0.010
polyBERT (Kuenneth and Ramprasad, 2023)	0.553 ± 0.011	0.759 ± 0.042	0.363 ± 0.037	0.526 ± 0.068	18.437 ± 0.560	0.618 ± 0.049	0.113 ± 0.003	0.172 ± 0.016
Transpolymer (Xu et al., 2023a)	0.453 ± 0.007	0.576 ± 0.021	0.326 ± 0.040	0.397 ± 0.061	17.740 ± 0.732	0.547 ± 0.051	0.096 ± 0.016	0.147 ± 0.093
MMPolymer (Wang et al., 2024)	0.431 ± 0.017	<u>0.503</u> ± 0.038	<b>0.286</b> ± 0.029	<b>0.390</b> ± 0.057	<b>16.814</b> ± 0.867	<u>0.511</u> ± 0.035	<b>0.087</b> ± 0.034	<u>0.061</u> ± 0.010
Uni-3DAR	<b>0.426</b> ± 0.022	<b>0.498</b> ± 0.048	<u>0.291</u> ± 0.022	<u>0.396</u> ± 0.072	<u>17.16</u> ± 1.498	<b>0.487</b> ± 0.034	<b>0.087</b> ± 0.011	0.066 ± 0.031

## C.7 POLYMER PROPERTY PREDICTION VIA PRETRAINING

Polymers, synthesized through various polymerization methods such as addition, ring-opening, and condensation, consist of repeating monomer units. These materials play essential roles across multiple fields, including materials science, drug design, and bioinformatics, necessitating accurate property prediction methods. Here, we demonstrate Uni-3DAR’s structure-level understanding capability by focusing on homopolymer property prediction.

**Dataset and Metric** Following prior research (Zhang et al., 2023b; Wang et al., 2024), we use eight publicly available polymer property datasets (Egc, Egb, Eea, Ei, Xc, EPS, Nc, and Eat), obtained via density functional theory (DFT) calculations. Given that all tasks involve structure-level regression, we employ a robust evaluation strategy using 5-fold cross-validation with random splits, consistent with previous work (Wang et al., 2024). Results are reported as the root mean squared error (RMSE), averaged across three different random seeds.

**Baselines and Implementation** Baseline methods include ChemBERTa (Chithrananda et al., 2020), Uni-Mol (Zhou et al., 2023b), SML (Zhang et al., 2023b), PML (Zhang et al., 2023b), polyBERT (Kuenneth and Ramprasad, 2023), Transpolymer (Xu et al., 2023a), and MMPolymer (Wang et al., 2024). For pretraining, we represent homopolymers as specialized molecular structures using the star substitution strategy proposed in (Wang et al., 2024). The model is pretrained using our masked next-token prediction strategy for 1 million steps with a batch size of 128. All other experimental details follow the settings previously described in the molecular property prediction experiments.

During fine-tuning, we adopt structure-level understanding strategy with masked next-token prediction auxiliary generative loss. Training is capped at 200 epochs. We thoroughly investigate various hyperparameter combinations by using three different batch sizes (32, 64, 128) and three learning rates (5e-4, 1e-4, 3e-4), creating nine unique configurations. Each configuration is tested by training models three times with different random seeds. We follow the same 5-fold split index align with (Wang et al., 2024), by averaging the best validation metrics in each fold. Subsequently, we present the mean performance along with the standard deviation across three seeds.

## D MORE EXPERIMENTS

### D.1 MUTUAL BENEFITS OF GENERATION AND UNDERSTANDING TASKS

In the previous experiments, we applied Uni-3DAR independently to each task to ensure fair comparisons with established approaches, rather than employing joint training across multiple tasks and diverse data sources. Although earlier results already demonstrate Uni-3DAR’s effectiveness, the

1566 Table 10: In the molecular pretrained representation task, incorporating a generation loss during downstream  
 1567 fine-tuning improves performance.

	HOMO ↓ (Hartree)	LUMO ↓ (Hartree)	E1-CC2 ↓ (eV)	E2-CC2 ↓ (eV)
Uni-3DAR w/o Gen. loss	0.0052	0.0049	0.0063	0.0077
Uni-3DAR	<b>0.0048</b>	<b>0.0044</b>	<b>0.0056</b>	<b>0.0067</b>

1573 Table 11: In the QM9 unconditional generation task, incorporating a structure-level understanding task further  
 1574 enhances the quality of the generated samples.

	QM9			
	Atom Sta(%))↑	Mol Sta(%))↑	Valid(%))↑	V × U(%))↑
Uni-3DAR	99.4	93.7	98.0	<b>94.0</b>
Uni-3DAR w/ Structure Und. loss	<b>99.6</b>	<b>95.8</b>	<b>98.5</b>	93.1

1582 advantages of joint training, particularly combining generation and understanding tasks, remain less  
 1583 explored. Due to resource limitations, comprehensive large-scale joint training was not feasible in  
 1584 this paper. Nonetheless, this subsection presents two additional experiments that clearly illustrate  
 1585 how generation and understanding tasks can mutually reinforce each other, highlighting the potential  
 1586 for enhanced performance through joint training in Uni-3DAR.

1587 The first experiment leverages the pretrained molecular representation described in Sec. 3.6. Typically,  
 1588 during downstream fine-tuning, we include an auxiliary generative loss by predicting ground-  
 1589 truth atom types and positions with the proposed masked next-token prediction. To investigate the  
 1590 contribution of this auxiliary generation task, we performed an ablation experiment by removing the  
 1591 generation loss during fine-tuning (results shown in Table 10). The results indicate a notable perfor-  
 1592 mance drop without the generation loss, clearly demonstrating that generative training significantly  
 1593 strengthens structure-level understanding.

1594 The second experiment builds upon the unconditional 3D molecule generation task using the QM9  
 1595 dataset described in Sec. 3.1. Previously, to align with prior studies, we used only 3D molecular  
 1596 structure data. Here, we additionally incorporate a structure-level understanding task by predicting  
 1597 the molecular property  $U$  (internal energy at 298.15 K), with results shown in Table 11. Models  
 1598 trained with this auxiliary structure-level understanding task consistently outperform those without,  
 1599 especially in metrics such as molecular stability and validity. This demonstrates that structure-level  
 1600 understanding significantly enhances generative performance.

1601 In summary, these experiments robustly illustrate that generation and understanding tasks positively  
 1602 reinforce one another. The findings underscore that integrating diverse datasets and joint task training  
 1603 can establish a more powerful and effective foundation model for 3D structural modeling.

## 1604 D.2 ABLATION STUDY

1606 We conducted comprehensive ablation experiments on QM9 generation task (Sec. 3.1) to eval-  
 1607 uate the contributions of key components in Uni-3DAR. The experimental results, summarized in  
 1608 Table 12, lead to the following insights:

- 1610 **1. Masked Next-Token Prediction significantly enhances generation performance.** In exper-  
 1611 iment No.2, we followed previous work (Ibing et al., 2023) that merely appends the position of the  
 1612 next token to the current token, without using our proposed masked next-token prediction. Com-  
 1613 paring experiments No.1 and No.2 clearly demonstrates that our proposed masked next-token  
 1614 prediction substantially outperforms this baseline approach.
- 1615 **2. 2-Level Subtree Compression boosts efficiency without compromising performance.** Exper-  
 1616 iment No.3 evaluates performance without 2-level subtree compression. Comparing No.1 (with  
 1617 compression) and No.3 (without compression), we observe that using subtree compression re-  
 1618 duces token count by approximately 6x, leading to significantly faster training with comparable  
 1619 results. Interestingly, experiment No.4 (No.2 without subtree compression) outperforms No.2.  
 This indicates that while subtree compression alone may slightly impact performance negatively

1620 Table 12: Ablation Studies for Uni-3DAR. MNTP (Masked Next-Token Prediction) boosts performance, while  
 1621 2LSC (2-Level Subtree Compression) enhances efficiency. Uni-3DAR integrates both techniques to balance  
 1622 effectiveness and efficiency. Token-level diffusion loss (diff. loss) performs comparably to our proposed simple  
 1623 autoregressive head. Training cost is measured using 4 NVIDIA 4090 GPUs.

No.	Settings	QM9				# AVG. Tokens	Training Cost ↓
		Atom Sta(%)↑	Mol Sta(%)↑	Valid(%)↑	$V \times U(%)↑$		
1	Uni-3DAR	99.4	93.7	98.0	94.0	160	6.9h
2	1 w/o MNTP	98.7	88.2	97.0	91.5	80	6h
3	1 w/o 2LSC	99.4	94.4	98.2	92.1	1060	20h
4	2 w/o 2LSC	99.3	94.2	97.7	92.7	530	11h
5	2 w/o octree	87.7	25.3	72.1	65.7	18	3h
6	1 w/ diff. loss	99.4	93.6	98.2	94.0	160	7.8h
7	5 w/ diff. loss	88.3	35.4	67.3	46.5	18	5.6h

1633  
 1634 (No.2 vs. No.4), when combined with masked next-token prediction (No.1 vs. No.3), it achieves  
 1635 comparable performance efficiently.

1636

3. **Coarse-to-fine octree tokens provide essential spatial information.** In experiment No.5, we  
 1637 removed octree tokens, significantly degrading model performance. Without coarse-to-fine to-  
 1638 kenization, the model degrades to atom-based autoregressive prediction of both atom types and  
 1639 positions, a much more challenging task. Our coarse-to-fine octree tokenization method effec-  
 1640 tively provides positional priors from preceding levels, substantially enhancing performance. This  
 1641 clearly validates the importance of the coarse-to-fine tokenization strategy for 3D structural gen-  
 1642 eration.
4. **Token-level diffusion loss yields comparable performance to the autoregressive head but  
 1643 with lower efficiency.** Our default generation head uses a simple autoregressive head (refer to  
 1644 Alg. 1) to sequentially predict atom types and in-cell positions. We examined whether employing  
 1645 a more powerful head, such as the token-level diffusion loss from MAR (Li et al., 2024b), could  
 1646 further enhance performance. Experiment No.6, utilizing the diffusion head, achieved similar  
 1647 results but required more computational time. Therefore, we opt for the simpler, more efficient  
 1648 autoregressive head by default.
5. **Combining atom-based autoregressive and diffusion losses without spatial tokenization is in-  
 1649 sufficient.** Recent works have explored improving atom-based autoregressive generation through  
 1650 token-level diffusion losses (Zhang et al., 2025). We tested this approach by adding a token-  
 1651 level diffusion loss to experiment No.5, resulting in experiment No.7. Although No.7 performed  
 1652 slightly better than No.5, it remained significantly inferior to the proposed Uni-3DAR. This  
 1653 underscores that comprehensive spatial information, as provided by our tokenization strategy, is  
 1654 crucial, mere integration of diffusion-based methods into atom-based autoregressive model, with-  
 1655 out spatial tokenization, cannot achieve substantial performance improvements.

### 1658 D.3 INFERENCE SPEED

1659 We benchmarked Uni-3DAR against the diffusion-based generative model GeoLDM (Xu et al.,  
 1660 2023b) on QM9 generation task (Sec.3.1) by evaluating the throughput (i.e., the number of  
 1661 molecules generated per second). Model throughput was evaluated across a range of batch sizes,  
 1662 with all experiments conducted on a single Nvidia 4090 GPU. As shown in Fig.4, Uni-3DAR con-  
 1663 stantly outperforms the diffusion-based approach in sampling efficiency, achieving significantly  
 1664 reduced generation times across all tested settings. In particular, at larger batch sizes, Uni-3DAR  
 1665 is approximately 21.8x faster than GeoLDM, and even at a small batch size of 64, it remains about  
 1666 7.5x faster. Additionally, we assessed the inference overhead introduced by masked next-token pre-  
 1667 diction. Thanks to our optimizations (Sec. B), we find that masked next-token prediction incurs  
 1668 only a 15% to 30% slowdown. Given its substantial performance gains, this additional cost is well  
 1669 justified.

1670  
 1671  
 1672  
 1673

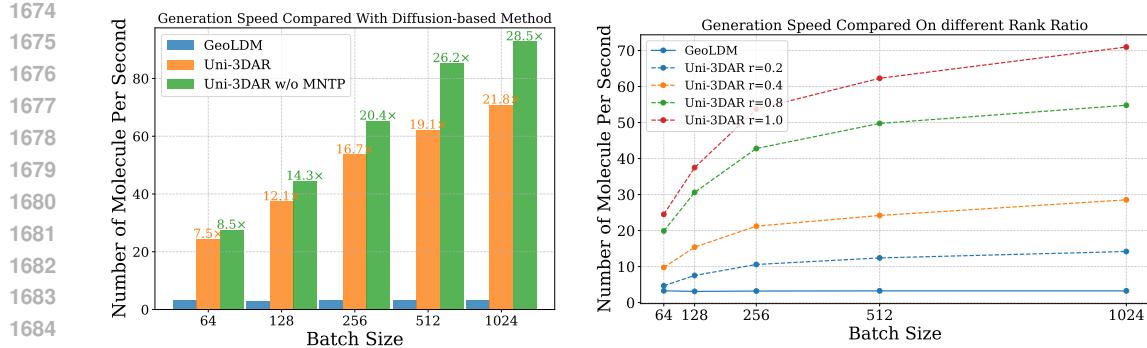


Figure 4: **Left:** Uni-3DAR generation speed on different batch sizes compared with the diffusion-based method; **Right:** Uni-3DAR generation speed on different rank ratios  $r$  compared with the diffusion-based method (higher is better).

## E ILLUSTRATION OF THE GENERATED EXAMPLES

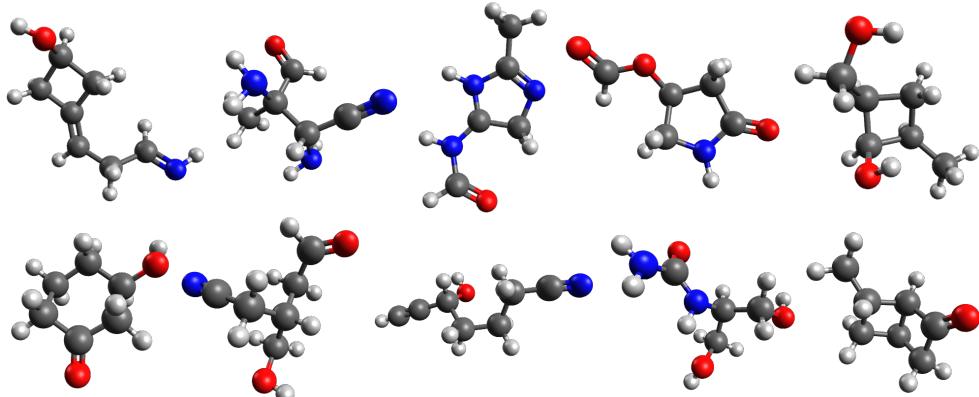


Figure SI-1: Unconditional 3D molecular generation samples of QM9 dataset.

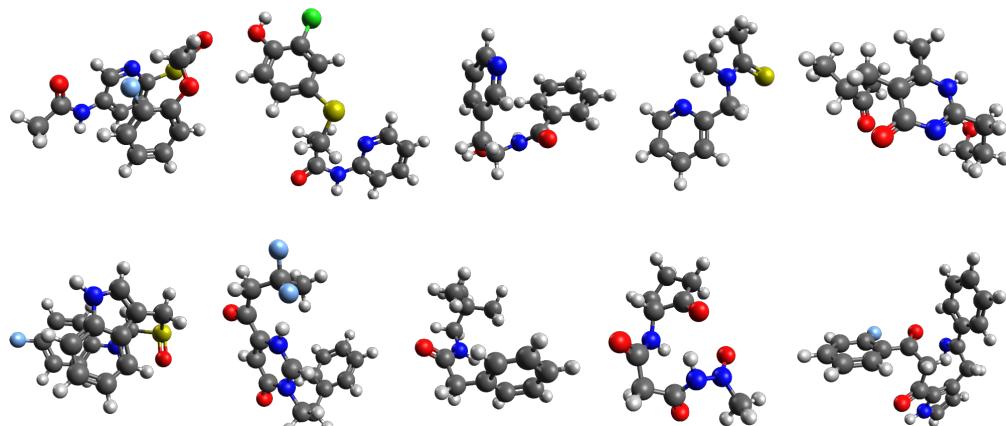


Figure SI-2: Unconditional 3D molecular generation samples of GEOM-DRUG dataset.

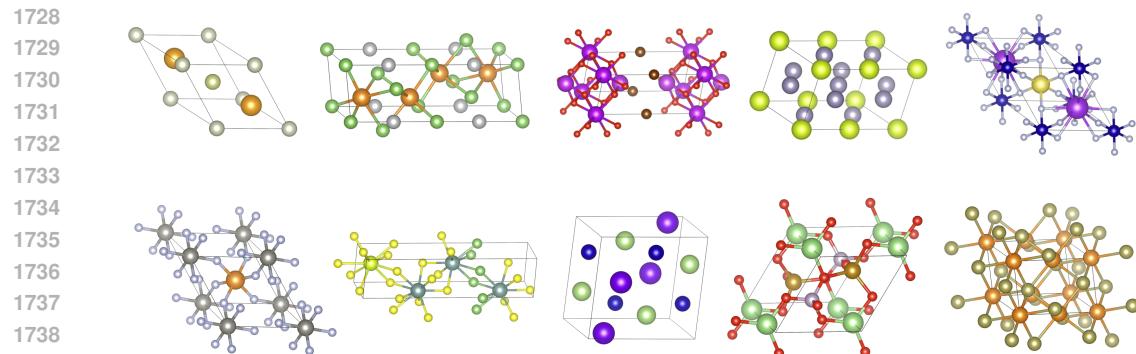


Figure SI-3: De novo crystal generation samples of MP-20 dataset.

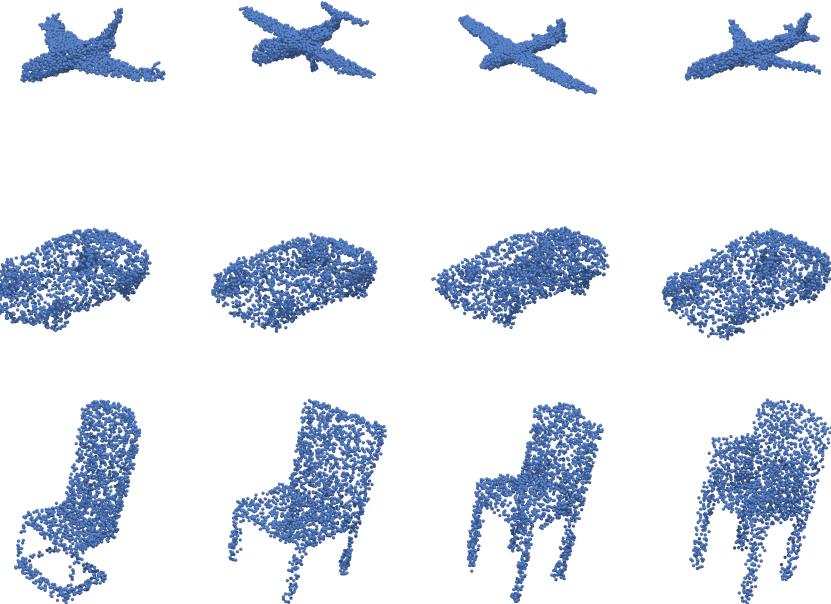


Figure SI-4: Macroscopic 3D object generation samples of ShapeNet dataset.

---

**F LLM USAGE DETAILS**

During the preparation of this manuscript, we utilized Large Language Models (LLMs), Google’s Gemini-2.5-pro, for assistance. The use of these models was strictly limited to improving the language and readability of the text. Specific applications included proofreading for grammatical errors, refining sentence structure for clarity, and ensuring a consistent and professional tone throughout the paper. The core scientific ideas, methodologies, experimental results, and conclusions presented in this work were conceived and articulated entirely by the human authors. All AI-generated suggestions were carefully reviewed and edited by the authors to ensure that the final text accurately reflects our original research and intent. The authors take full responsibility for the scientific content and integrity of this paper.

1782  
17831784  
1785  
1786  
1787  
1788  
1789  
1790  
1791  
1792

1793

1794

1795

1796

1797

1798

1799

1800

1801

1802

1803

1804

1805

1806

1807

1808

1809

1810

1811

1812

1813

1814

1815

1816

1817

1818

1819

1820

1821

1822

1823

1824

1825

1826

1827

1828

1829

1830

1831

1832

1833

1834

1835