GRAPH DIFFUSION TRANSFORMERS ARE IN-CONTEXT MOLECULAR DESIGNERS

Gang Liu[1], **Jie Chen**[2], **Yihan Zhu**[1], **Michael Sun**[3]
Tengfei Luo[1], **Nitesh V. Chawla**[1], **Meng Jiang**[1]
[1]University of Notre Dame [2]MIT-IBM Watson AI Lab, IBM Research [3]MIT CSAIL
{gliu7, mjiang2}@nd.edu

ABSTRACT

In-context learning lets large models adapt to new tasks from a few demonstrations, but it has shown limited success in molecular design, where labeled data are scarce and properties span millions of biological assays and material measurements. We introduce demonstration-conditioned diffusion models (DemoDiff), which define task contexts through molecule–score examples instead of texts. These demonstrations guide a denoising Transformer to generate molecules aligned with target properties. For scalable pretraining, we develop a new molecular tokenizer with Node Pair Encoding that represents molecules at the motif level, requiring 5.5× fewer nodes. We pretrain a 0.7B parameter model on datasets covering drugs and materials. Across 33 design tasks in six categories, DemoDiff matches or surpasses language models 100–1000× larger and achieves an average rank of 4.10 compared to 6.56–17.95 for 19 baselines. These results position DemoDiff as a molecular foundation model for in-context molecular design.

1 INTRODUCTION

In-context learning (ICL) is the emergent capability of large models to infer task-specific concepts from a few demonstrations (Xie et al., 2021). ICL has been studied in large language models (LLMs), but was found less effective for molecular design than specialized methods (Liu et al., 2024b). These specialized models often depend on extensive Oracle calls (Gao et al., 2022) or large labeled datasets beyond what context examples provide. Molecular tasks, however, involve millions of types, many with only a few labeled examples (Zdrazil et al., 2024). Such examples are enough to form task contexts but insufficient to train a new model. This trade-off motivates our *in-context molecular design*, which combines the flexibility of ICL with the efficiency of molecular domain knowledge.

Molecular structures and properties are discrete graphs and numbers with varying scales and units. Directly adapting the autoregressive framework from LLMs is infeasible (Brown et al., 2020) for in-context molecular designs, as the input and output of language data are text in sequential order. Diffusion models show promise for molecular structures (Vignac et al., 2022), and Graph Diffusion Transformers (Graph DiTs) are effective for modeling their joint distribution with properties (Liu et al., 2024c). However, Graph DiTs have been studied with at most five properties represented in a single vector. In practice, molecular properties span millions of assays in biology, including functions, binding, ADME, and toxicity, as well as material properties such as gas permeability, thermal conductivity, and glass transition temperature (Figure 1). Representing all properties in one-hot vector with millions of dimensions is inefficient, produces sparse pretraining data since many assays have fewer than ten labels, and limits generalization to unseen properties in downstream tasks.

Instead of a property vector with a large embedding table, we use demonstrations to define the task context for molecular design. As shown in Figure 1, the demonstrations consist of a set of molecules with scores in $[0, 1]$ and molecular design is framed as a query for the target score of 1. Molecules in the context do not follow a strict order, and their scores serve as *relative* positions to the target, functioning as a replacement for position IDs in Transformers. (Xie et al., 2021) described ICL in LLMs as implicit Bayesian inference over latent *concepts* expressed by examples in the prompt. Similarly, each task in Figure 1 shares the concepts defined by the joint distribution of molecules and their scores. The denoising Transformer in the Graph DiT attends to the context, implicitly

Figure 1: In-context molecular design with DemoDiff. Each demo is de ned as a score–molecule pair, and a set of them forms the task context as conditions. After pretraining on large and diverse tasks, DemoDiff serves as a foundation model for designing molecules in new task contexts. Scores represent relative distances to the target and are converted from raw labels, as shown in Section 3.3.

extracts concepts, and uses them to guide the reverse process to re ne the structure. An example generation trajectory is shown in Figure 7. A simple way to represent the task concept is to use positive demonstrations, such as molecules close to the target or active in the assay. However, positive examples alone may be insuf cient, as they can overlap across tasks due to factors such as task relatedness (e.g., activity in non-small-cell lung cancer but across different cell lines) or due to sampling bias when the set of positives is extremely sparse (e.g., only one positive example shared by two tasks). To address this, we form the task context using not only positive but also medium and negative examples, providing a more complete representation of the task concepts.

With these task contexts, we propose demonstration-conditioned diffusion models (DemoDiff) and pretrain a 0.7B model with a Graph DiT as the backbone, using over 140 H100 GPU days. To support ef cient pretraining, we introduce a molecular tokenizer trained with Node Pair Encoding (NPE) for motif-level representation. On average, it reduces the number of nodes by 5.5× compared to atom-level representations (Figures 9 and 10). The tokenizer iteratively merges neighboring nodes and selects frequent motifs to construct vocabularies. Motifs are connected by directed edges that preserve bond types and attachment rules, ensuring lossless reconstruction. Graph DiTs naturally use this motif-level representation and attend to motif semantics for denoising. For pretraining, we construct a dataset of over 1.6 million tasks from 155K unique properties and one million molecules. It combines ChEMBL for drugs (Zdrazil et al., 2024) and multiple polymer data sources for materials (Otsuka et al., 2011; Thornton et al., 2012; Kuenneth et al., 2021). For ICL, we propose a consistency score as a con dence measure of whether a generation aligns more closely with higher-scoring molecules in demonstrations, effectively ltering out false positives in generation (Section 4.3).

We evaluate DemoDiff on 33 design tasks across six categories. With 0.7B parameters, it matches or surpasses LLMs 100–1000× larger in generating diverse, high-scoring molecules. Compared to 19 baseline models (average ranks 6.56–17.95), DemoDiff ranks 4.10, demonstrating its strength as a molecular foundation model. The new molecular tokenizer further improves representation ef ciency.

2 PRELIMINARIES

2.1 IN-CONTEXT LEARNING WITH DEMONSTRATIONS

In ICL, the context is a set of demonstrations $C = \{e_i\}_{i=1}^{L}$, where each e_i is an input–label pair. Following (Xie et al., 2021), we assume C re ects a latent concept θ from a family of concepts

. For example, for a paragraph about Albert Einstein, the latent concept may be biography. Given a query Q, a foundation model with ICL generates an output X by marginalizing over :

$$p(X \mid C, Q) = \int p(X \mid , C, Q)\, p(\mid C, Q)\, d \; : \tag{1}$$

Here (C, Q) form the prompt. If $p(\mid C, Q)$ concentrates on the prompt concept with more demonstrations, then the model identifies and applies that concept through marginalization. ICL can thus be implicit Bayesian inference. All context, query, and outcomes are texts in language modeling. In inverse molecular design, we have molecule-score pairs as demonstrations. They capture the latent task concept, with semantics like the task descriptions in Figure 1. Q is the target score, and X is the molecule to be designed. We focus on a new molecular foundation model for ICL.

2.2 Molecular Design with Graph Diffusion Transformers

Molecules are discrete graphs $X = (A, B)$, where A denotes the set of atoms and B the set of bonds. These structures are commonly modeled using discrete diffusion processes (Vignac et al., 2022; Liu et al., 2024c). Graph DiTs concatenate atom and bond features into the input format of standard Transformers. Given X, for each atom $a_i \in A$ with d_i neighbors, Graph DiTs define a token as $x = \{a_i, \{b_{ij}\}_{j=1}^{d_i}\}$, where $b_{ij} \in B$ encodes the bond type (single, double, triple, or none). Each token is represented by a feature vector $x \in \mathbb{R}^F$, formed by concatenating the one-hot encoding of the atom type and the connection types to all other atoms (either a bond type or a null type indicating no connection). Discrete diffusion has a transition matrix Q, initialized based on the frequency of atoms and bonds in the training set. At step t, $[Q^t]_{ij} = q(x_j^t \mid x_i^{t-1})$ for $i, j \in [1, F]$.

The forward diffusion with Q is: $q(x^t \mid x^{t-1}) = \mathrm{Cat}(x_t; p = x^{t-1}Q^t)$, where $\mathrm{Cat}(x; p)$ denotes two separate categorical sampling for atoms and bonds with probabilities from p. Starting from the original data $x = x^0$, we have $q(x^t \mid x^0) = \mathrm{Cat}(x^t; p = x^0\bar{Q}^t)$, where $\bar{Q}^t = \prod_{i=1}^{t} Q^i$. The forward diffusion gradually corrupts data points. When the total time step T is large enough, $q(x^T)$ converges to a stationary distribution. The reverse process samples from $q(x^T)$ and gradually removes noise. The posterior distribution $q(x^{t-1} \mid x^t)$ is calculated as $q(x^{t-1} \mid x^t; x^0) \propto x^t (Q^t)^\top \odot x^0 \bar{Q}^{t-1}$. Given multiple properties $\{c_i\}_{i=1}^n$, the denoising model approximates $p_\theta(x^{t-1} \mid x^t; x^0; \{c_i\}_{i=1}^n)$ under property conditions. This model is trained by minimizing the negative log-likelihood of x^0 for

$$L_{DM} = \mathbb{E}_{q(x^0)} \mathbb{E}_{q(x^t \mid x^0)} \left[-\log p_\theta\left(x^0 \mid x^t; c_1, c_2, \dots, c_n\right) \right]; \tag{2}$$

In Graph DiTs (Liu et al., 2024c), the constraints $\{c_i\}_{i=1}^n$ are numerical or categorical property values. In this work, we explore them as demonstrations for in-context learning.

3 Learning Diffusion Model with Demonstrations

Figure 2 shows the generation process of DemoDiff, combining motif-based representation (Section 3.1) with graph diffusion transformers for in-context molecular generation (Section 3.2).

3.1 Molecular Graph Tokenization with Node Pair Encoding

More demonstrations help capture the latent task concept and are empirically useful (Bertsch et al., 2024). However, prior work (Liu et al., 2024c) uses atom-level molecular representations, similar to modeling text at the character level, which fundamentally limits the number of examples in context. For efficient representation, we merge frequent sub-molecular patterns into motifs $m = (\tilde{A}, \tilde{B}) \subseteq X$, where $\tilde{A} \subseteq A$ and $\tilde{B} \subseteq B$ define a connected substructure. A molecule becomes a collection of disjoint motifs $M = \{m_i\}_{i=1}^n$ such that: (1) $\tilde{A}_i \cap \tilde{A}_j = \emptyset$ for all $i \neq j$; (2) $\bigcup_{i=1}^n \tilde{A}_i = A$. Then, each edge $e_{ij} \in E$ is directed from a source motif m_i to a target motif m_j, with two associated attributes: (1) bond type, and (2) attachment index, indicating the atom within m_i from which the bond originates. This abstraction induces a tokenizer with two functions. They are $\mathrm{tokenizer.encode}: X = (A, B) \to \hat{X} = (M, E)$ that compresses the atom-level graph into motif-level form and $\mathrm{tokenizer.decode}: \hat{X} = (M, E) \to X = (A, B)$ for reconstruction. The tokenizer uses two vocabularies M for motif types and E for edge types. It starts with the 118 atom

Figure 2: Demonstration-conditioned diffusion generation. In the reverse process, DemoDiff starts from random noise and denoises molecules conditioned on molecule–score demonstration pairs. A tokenizer encodes each molecule into a motif-level representation. Motifs are connected by edges with two attributes shown in square brackets: (1) bond type and (2) attachment index. The input to the transformers consists of tokens from the denoised molecules and the context of demo tokens.

types from the periodic table and one "*" for the polymerization point, which form the initial motifs in $\mathcal{M}$. This guarantees that, in the worst case, a new molecule can still be represented at the atom level. In each iteration, the tokenizer merges the most frequent neighbors until no further merge is found in $\mathcal{M}$, then proceeds with the corresponding connections between motifs.

To construct $\mathcal{M}$, existing methods, such as BRICS (Degen et al., 2008) or molecular grammars (Sun et al., 2025; 2024), rely on domain-specific heuristics based on chemical reactions or expert knowledge. The resulting vocabularies are independent of the pretraining data, often missing frequent motifs. To address this limitation, we propose Node Pair Encoding (NPE), a frequency-based algorithm for molecular graphs inspired by BPE, as outlined in Algorithm 1. We initialize $\mathcal{M}$ with elements from the periodic table and polymerization points "*". NPE iteratively performs three steps: (1) Neighborhood merge For each molecule $X \in \mathcal{D}$ in the dataset $\mathcal{D}$ and current $\mathcal{M}$, we identify candidate motifs by merging adjacent substructures that appear in X. (2) Frequency selection: The most frequent candidate motif is selected and added to $\mathcal{M}$. (3) Graph update: Each $X \in \mathcal{D}$ is updated by replacing instances of merged motif pairs with the new motif.

Constrained NPE: The standard NPE may produce multiple directed edges from a motif m_i to another m_j when decomposing ring structures (e.g., aromatic rings). It leads to ambiguity during decoding since e_{ij} does not uniquely determine the attachment index within m_j. To address this, we introduce constraints such as rings into NPE at two stages. During initialization, we traverse each molecule to identify its set of maximal connected rings, denoted $\mathcal{R}$, compute their frequencies, and include the top-K_{ring} most frequent rings in the initial vocabulary $\mathcal{M}$. During motif merging, any $m \in \mathcal{M}$ is merged with a ring $r \in \mathcal{R} \setminus \mathcal{M}$ as a complete unit, rather than merging individual atoms within r. This strategy integrates frequent rings into the vocabulary while preserving atom-level representations for rare rings, avoiding reconstruction ambiguity. We compare NPE with fingerprints (Rogers & Hahn, 2010) and virtual nodes (Hwang et al., 2022) in appendix A.

Using NPE, a molecule is represented as tokens $\{x_i\}_{i=1}^{n}$, where $x_i = \{m_i; \{e_{ij}\}_{j=1}^{d_i}\}$, with m_i denoting a motif and e_{ij} the associated edges. We set the motif vocabulary size to 3000 ($K_{\text{ring}} = 300$), with details and analysis provided in appendix C.2. As shown in Figure 2, an example input with 38 atoms can be compactly expressed using four motifs. An empirical comparison of atom- and motif-level representations over 1 million pretraining molecules is given in Figure 3b, with an average compression ratio of $44.5/6.446 \approx 2.569$, reducing the median count from 30 atoms to 5 motifs.

3.2 In-Context Learning with Graph Diffusion Transformers

We construct the dataset $\mathcal{D}_{\text{pt}} = \{(C_i; Q_i; X_i)\}_{i=1}^{N_{\text{pretrain}}}$ for pretraining, where each task consists of a context of molecule–score pairs C_i, a query score Q_i, and a target molecule X_i. To pretrain DemoDiff

for in-context inverse molecular design, we replace the property conditions in Eq. (2) with C and Q:

$$L_{\text{pretrain}} = E_{q(x^0)} E_{q(x^t|x^0)} \quad \log p \quad x^0 \mid x^t; C; Q \quad : \tag{3}$$

With large and diverse pretraining data and scalable Transformers (Peebles & Xie, 2023), DemoDiff learns to infer the latent task concept to generate the target molecule and can serve as a foundation model for ICL. In a task, it performs implicit Bayesian inference over diffusion trajectories.

ICL with Context Consistency (appendix B.4): Given a query Y, demonstrations S_i are divided into positive, medium, and negative groups. For a generated molecule X, we compare its ngerprint-based similarity with these groups to assess whether it follows the relation pos $>$ med $>$ neg. This yields a consistency score that measures how well the generated molecule aligns with the relative relations in the context. In experiments (Section 4.1), we use this score to select high consistent generations before conducting the nal evaluation with Oracles.

3.3 Model Design and Pretraining

(a) Property rank vs. frequency in $\mathcal{D}$. (b) Node counts (50; full in Figure 9).

Figure 3: Pretraining data statistics for property rank-frequency and node count density.

Model Designs: Figure 2 illustrates the model architecture. For the i-th task in the pretraining set $(C_i; X_i; Y_i) \in \mathcal{D}$, the property score Y_i is scaled within $[0; 1]$, providing positional signals for both demonstration molecules and the target. These scalar values are encoded using Rotary Position Embedding (RoPE) (Su et al., 2024). DemoDiff uses a tokenizer to process the atom-level representation and de nes a maximum context length for the number of motif tokens. The context includes the target along with as many demonstration tokens as t within the target length. Since molecules in the context are structurally disjoint, edge connectivity implicitly delineates context boundaries, removing the need for explicit delimiter tokens. Details are in appendix B.

Pretraining: To construct the pretraining dataset as illustrated in Section 3.2, we use the ChEMBL database (Zdrazil et al., 2024), the largest collection of biological assays, containing over 2.5 million molecules and 1.7 million assay records. To increase chemical diversity for materials discovery, we augment ChEMBL with polymer datasets from multiple sources (Liu et al., 2024b; Kuenneth et al., 2021), including properties such as thermal conductivity, free volume fraction, and glass transition temperature. For biological assays, we generate tasks by selecting a molecule-activity pair as the target and treating other molecule-activity pairs as context. The target is assigned a score of 1, and context scores are computed by normalizing differences in ChEMBL values (negative log of bioactivity measures such as IC50 and potency) to the interval $[0; 1]$. We restrict targets to bioactive molecules with pChEMBL > 6. For polymers, we apply the same strategy: each polymer is used as a target, and its property value is normalized against those of other polymers to form context-target pairs. We partition context examples into three groups by normalized scores: positive $[0:75; 1]$, medium $[0:5; 0:75)$, and negative $[0; 0:5)$, with up to 15 demonstrations from each. The nal dataset comprises around 1 million molecules with 155K unique assays or properties, yielding 1.6 million tasks. As shown in Figure 3a, the frequency distribution of assays and properties follows Zipf's law, $P(Y_{\text{rank}}) \propto \text{rank}^{-1:13}$, consistent with patterns in language corpora. These 1 million molecules are used to initialize the motif vocabulary via NPE, which provides a more compact representation by reducing node counts (Figure 3b). We further extract edge connections to construct an edge

Table 1: We compute the oracle and diversity scores from the Top-10 generations and report the harmonic mean of these two values. We group 33 tasks into six categories and report the mean within each category. The best results in each column are bolded. Task-specific results and additional metrics are provided in appendix D.2.

Task	Drug Rediscovery	Drug MPO	Structure Constrained	Drug Design	Target Based	Material Design	Avg Rank	Total Sum
# Tasks	7	7	5	4	5	5	33	33
Molecular Optimization Methods with 100 Oracle Calls								
GraphGA	0.36±0.07	0.52±0.19	0.43±0.21	0.41±0.32	0.76±0.04	0.58±0.11	6.56	16.65
REINVENT	0.37±0.08	0.52±0.17	0.43±0.21	0.42±0.32	0.76±0.03	0.00±0.00	8.44	13.84
GPBO	0.37±0.07	0.51±0.18	0.42±0.22	0.39±0.33	0.76±0.03	0.60±0.21	6.89	16.65
STONED	0.36±0.07	0.52±0.19	0.39±0.24	0.37±0.36	0.76±0.04	0.00±0.00 (4 pSMILES)	8.18	13.42
Genetic GFN	0.36±0.10	0.51±0.22	0.38±0.29	0.33±0.37	0.76±0.03	0.00±0.00 (3 pSMILES)	9.28	13.04
GenMol	0.42±0.08	0.51±0.17	0.42±0.25	0.55±0.21	0.69±0.03	0.62±0.08	7.98	17.36
Molecular Optimization Methods with 10000 Predictor Calls								
GraphGA	0.37±0.09	0.50±0.18	0.45±0.29	0.49±0.26	0.64±0.06	0.55±0.14	9.30	16.30
REINVENT	0.30±0.13	0.23±0.23	0.25±0.24	0.38±0.26	0.17±0.13	0.51±0.16	13.60	9.82
GPBO	0.33±0.10	0.45±0.22	0.43±0.29	0.49±0.15	0.74±0.03	0.42±0.24	9.87	15.35
STONED	0.33±0.10	0.40±0.20	0.50±0.28	0.27±0.10	0.20±0.27	0.00±0.00 (6 pSMILES)	12.01	9.66
Genetic GFN	0.14±0.08	0.08±0.16	0.08±0.11	0.12±0.11	0.15±0.15	0.00±0.00 (5 pSMILES)	17.95	3.14
GenMol	0.34±0.07	0.28±0.23	0.21±0.29	0.39±0.19	0.52±0.05	0.21±0.23	12.63	10.37
Conditional Generation Models								
LSTM	0.39±0.25	0.16±0.07	0.55±0.32	0.33±0.35	0.72±0.04	0.16±0.11	11.77	12.30
Graph-DiT	0.43±0.21	0.50±0.18	0.58±0.34	0.48±0.37	0.71±0.04	0.55±0.17	8.53	17.64
Learning from In-Context Demonstrations								
DeepSeek-V3	0.45±0.18	0.51±0.20	0.49±0.24	0.65±0.18	0.64±0.06	0.39±0.24	8.08	16.90
GPT-4o	0.47±0.21	0.53±0.20	0.52±0.30	0.48±0.40	0.73±0.05	0.43±0.16	7.89	17.25
Qwen-Max	0.15±0.21	0.17±0.15	0.32±0.32	0.29±0.29	0.19±0.26	0.10±0.18	15.39	6.46
Llama3.1-8B-FT	0.21±0.13	0.24±0.23	0.24±0.34	0.31±0.31	0.02±0.05	0.29±0.40	14.91	6.04
Qwen3-8B-FT	0.37±0.19	0.27±0.16	0.26±0.34	0.46±0.22	0.67±0.07	0.44±0.39	10.96	12.23
DemoDiff (Ours)	0.44±0.21	0.54±0.23	0.56±0.33	0.79±0.11	0.78±0.05	0.67±0.11	4.10	20.10

vocabulary capturing motif-to-motif connectivity. Finally, we pretrain a DemoDiff model with 0.7B parameter on Eq. (3), using 146 H100 GPU days. Details are provided in appendix C.

4 EXPERIMENT

Setups: We curate 33 downstream tasks (see Table 1 and appendix D.1) across six categories to evaluate DemoDiff against 19 baselines. These tasks are primarily curated by domain experts and are distinct from pretraining. We include eight molecular optimization methods and two conditional generation models (LSTMs and Graph DiT (Liu et al., 2024c)), and LLMs (DeepSeek-V3, GPT-4o, and Qwen-Max). We select the top four molecular optimization algorithms from the PMO benchmark (Gao et al., 2022) (out of 25 evaluated methods) and two recent methods, Genetic GFN (Kim et al., 2024) and GenMol (Lee et al., 2025), under two settings: 100 oracle calls and 10,000 predictor calls. For evaluation, we generate 10 valid, unique, and novel molecules per task and score them with Oracles. We report the harmonic mean over two dimensions: (a) averaged oracle scores and (b) the diversity score Eq. (8). Each task has up to 450 molecule–score pairs, evenly divided into positive $[0.75, 1]$, medium $[0.5, 0.75)$, and negative $[0, 0.5)$ groups. Each task has an Oracle for evaluation. We use all molecules to train the task-specific predictor for predictor calls or to train conditional generation models directly. For LLMs with ICL, we include three closed-sourced models. We fine-tune Llama3.1-8B-FT and Qwen3-8B-FT on the DemoDiff pretraining set. Their demonstrations are randomly sampled with a similar budget for context. For the baselines, we follow their standard and original implementations to process the molecules.

4.1 PERFORMANCE ON DIVERSE MOLECULAR DESIGN TASKS

ICL achieves competitive performance with minimal supervision.Table 1 compares the harmonic mean of the top-10 generated molecules based on both task scores and structural diversity, while Table 5 (appendix) reports the top-1 scoring molecule per task. Under limited data and Oracle budgets, ICL methods perform comparably to, or better than, fully trained conditional generators and molecular optimization baselines. Among ICL-based methods, DeepSeek-V3, GPT-4o, and DemoDiff consistently attain top-tier average ranks. These ICL methods rely on tens of demonstrations per task, signi cantly fewer than the training data or Oracle calls required by other models or algorithms.

DemoDiff designs molecules with accurate scores and high diversity. Across six task categories, it performs best on property-driven tasks, including drug design with bioactivity targets, protein binding af nity, and material design for polymer gas separation. DemoDiff achieves the lowest average rank of 4.10, outperforming the best baseline, GraphGA (rank 6.56). ICL methods with LLMs produce high-scoring top designs (Table 5) but often generate structurally similar molecules. These do not necessarily align better with the target score while reducing diversity. In contrast, DemoDiff designs molecules with scores closer to the query and better structural diversity.

DemoDiff performs better on property-driven tasks than on structure-constrained ones.It scores 0.67–0.79 on drug and material design, but around 0.44–0.56 for rediscovery and structure-constrained tasks, where Oracle scoring is tied to the presence of speci c structures. While DemoDiff still ranks highly in structure-constrained tasks, its stronger results on property-driven tasks highlight its advantage in exploring chemical spaces with broader solution ranges.

4.2 Ablation Studies and Performance Analysis

Table 2: Performance across model sizes using harmonic mean scores from Top-10 generations

DemoDiff	Drug Rediscovery	Drug MPO	Structure Constrained	Drug Design	Target Based	Material Design
78M	0.39 ± 0.17	0.46 ± 0.24	0.47 ± 0.27	0.57 ± 0.31	0.73 ± 0.03	0.62 ± 0.13
311M	0.40 ± 0.17	0.46 ± 0.23	0.50 ± 0.28	0.53 ± 0.27	0.75 ± 0.04	0.62 ± 0.14
739M	0.44 ± 0.21	0.54 ± 0.23	0.56 ± 0.33	0.79 ± 0.11	0.78 ± 0.05	0.67 ± 0.11

Model Parameters: We pretrain DemoDiff with varying sizes: small (78.7M), medium (311M), and large (739M) parameters. Table 2 reports performance using the top-10 harmonic means of task score and diversity. We present averages with deviations across six categories. DemoDiff achieves reasonable scores even at small scale. At the medium scale, performance improves in most tasks except drug design, while the bene ts of parameter scaling become more evident (in ve out of six task categories) at larger scales

ICL with Demonstrations: Figure 4 studies two factors in demonstrations: (1) context length and (2) ratio of positive examples. In Figure 4a, longer context includes more molecular examples and supports better ICL performance. This aligns with the rationale of motif-level tokenization, which captures more examples within a xed context. Figure 4b shows that diverse demonstrations are important for ICL to represent the task accurately, while only positive examples are insuf cient. This is because positive, medium, and negative examples together provide a holistic view of the task context, and DemoDiff pretrained on such contexts is better able to infer latent concepts from diverse examples. In Figure 4b, we also observe that fewer positive examples may still yield reasonable results. We investigate this further in Section 4.3 to assess whether DemoDiff can infer positive examples (score > 0.5) using only negative examples with scores below 0.5.

ICL with Consistency Scores: We ablate consistency scores and analyze their correlation with target scores in Figure 5. Using the consistency score as a con dence lter improves performance across task categories, with gains from 0.8% to 27.5%. The second gure shows the correlation between the consistency and target scores. Moderate correlation appears in tasks with explicit structural constraints, such as drug rediscovery and structure-constrained design. For property-driven tasks (drug MPO and materials design), high ngerprint-based consistency with positive examples does not always correlate with high target scores. In these cases, latent concepts may rely on subtle substructures (e.g., methyl groups (Liu et al., 2022)) that standard ngerprints fail to capture. Interestingly, context consistency still improves performance in these tasks. A possible reason is that the score helps lter out false positive generations.

(a) Harmonic score vs. context length (b) Harmonic score vs. positive samples

Figure 4: Ablation studies on Albuterol drug rediscovery.

Figure 5: Ablation studies on context consistency scores: (1) left shows improvements; (2) right shows the relationship between consistency score and oracle scores.

4.3 CASE STUDIES

Figures 6 and 7 present two studies for DemoDiff. In extreme cases of inverse molecular design, demonstration sets may contain only negative examples, i.e., all scores < 0.5. In Figure 6, we study whether DemoDiff can still generate positive candidates when prompted solely with negative examples. Figure 6 presents the results for (a) structure-constrained design, (b) drug multi-objective optimization (MPO), and (c) target-based design. These findings suggest two insights: (1) negative demonstrations convey informative signals about the task concept, and (2) after pretraining, the posterior over the concept-to-structure mapping allows DemoDiff to generate desirable candidates that are aligned with the concept yet structurally distinct from the negative examples. Figure 7 is the generation trajectory from diffusion models. The task score, measured as structural similarity to Albuterol, rises from 0.22 at initial sampling to 0.74. This shows that the diffusion model refines the molecule toward the desired structure step by step with demonstrations.

5 RELATED WORK

Inverse Molecular Design: Molecular optimization uses diverse approaches, including genetic algorithms, Monte Carlo Tree Search (Jensen, 2019), and Bayesian optimization (Shahriari et al., 2015), applied to representations such as fingerprints, SMILES, graphs, and synthetic pathways (Gao et al., 2021). Gao et al. (2022) benchmarked 25 optimization methods and found that older models, such as genetic algorithms, remain competitive. However, existing benchmarks require on the order of 10,000 oracle calls, which is costly and limits applicability when single calls are expensive. Deep learning models offer an alternative by modeling the joint distribution of atoms and bonds without Oracle calls. GDSS applies noise and denoising in continuous space for graphs (Jo et al., 2022). DiGress (Vignac et al., 2022) introduces discrete noise through transition matrices based on marginal atom and bond distributions. Graph DiTs (Liu et al., 2024c) extend scalable diffusion transformers (Peebles & Xie, 2023) to discrete graphs. Yet, training diffusion models still requires hundreds of labeled molecules and is limited to specific tasks. Recent efforts explore chemical

Figure 6: Learning from negative demonstrations (score ≤ 0.5) to infer a target with score 1. All demonstrations are shown in Figures 14 to 16, with only three displayed here.

Figure 7: Diffusion trajectory for Albuterol drug rediscovery: we sample five intermediate diffusion steps and score them with the Albuterol Oracle, which computes similarity to the ground truth.

foundation models based on LLMs (Yu et al., 2024; Liu et al., 2024b), but their applications are either diverted to other molecular tasks, such as property prediction, or rely on fine-tuning within a restricted scope of design tasks.

In-Context Learning: ICL is an emergent ability observed in LLMs (Brown et al., 2020; Chan et al., 2022). Empirical and theoretical studies investigate this phenomenon from three perspectives: models, data, and learning mechanisms (Xie et al., 2021; Min et al., 2022). For the learning mechanism, ICL can be interpreted as implicit Bayesian inference (Xie et al., 2021), where pretraining data are generated from latent concepts and the posterior distribution marginalizes over them for inference. On the model side, Garg et al. (2022) trained Transformers from scratch on prompt-style input–label pairs of simple functions and found performance comparable to task-specific algorithms. Bhattamishra et al. (2023) compared Transformers with attention-free models and showed they do not match Transformer performance across tasks. On the data side, Chan et al. (2022) found that Transformers outperform recurrent models (e.g., LSTMs) on data with distributional properties resembling natural language, such as burstiness (words appearing in clusters) and query tasks with many rare classes. Singh et al. (2025) analyzed the strategy competition between ICL and in-weight learning, showing that the asymptotic strategy depends on in-weight information but is also context-constrained. This aligns with (Chan et al., 2022), suggesting that a foundation model should support both capacities. A skewed Zipfian distribution over tasks (e.g., Figure 3a) balances learning by storing common task information in weights while developing ICL ability from the long tail of rare tasks. Besides LLMs, ICL is also studied on graphs (Huang et al., 2023), but mainly for node and edge classification tasks.

6 CONCLUSION

We presented DemoDiff, a demonstration-conditioned diffusion Transformer model for in-context molecular design. We constructed a large-scale pretraining dataset with over one million molecules and 155K unique biological assays and material properties, yielding millions of demonstration–target pairs. Using this dataset, we pretrained a 0.7B-parameter model and showed that it matches or outperforms much larger LLMs and ranks higher than domain- and task-speci c methods. To support scalable pretraining, we introduce Node Pair Encoding, a motif-level graph tokenizer that ef ciently represents molecules with fewer nodes while preserving reconstruction. Experiments demonstrate that DemoDiff is a promising molecular foundation model, highlighting its potential to scale further with larger models, broader datasets, and greater compute.

ACKNOWLEDGMENTS

This work was partially supported by NSF IIS-2142827, IIS-2146761, IIS-2234058, and CBET-2332270. We also appreciate the support from the Foundation Models and Applications Lab of Lucy Institute, ND-IBM Tech Ethics Lab, and OpenAI Researcher Access Program.

REPRODUCIBILITY STATEMENT

We provide pretraining and inference code in the supplementary materials. The appendix describes the method and model settings to ensure reproducibility.

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023.

Amr Alhossary, Stephanus Daniel Handoko, Yuguang Mu, and Chee-Keong Kwoh. Fast, accurate, and reliable molecular docking with quickvina 2. Bioinformatics, 31(13):2214–2216, 2015.

Amanda Bertsch, Maor Ivgi, Emily Xiao, Uri Alon, Jonathan Berant, Matthew R Gormley, and Graham Neubig. In-context learning with long-context models: An in-depth exploration. arXiv preprint arXiv:2405.00200, 2024.

Satwik Bhattamishra, Arkil Patel, Phil Blunsom, and Varun Kanade. Understanding in-context learning in transformers and llms by learning to learn discrete functions. arXiv preprint arXiv:2310.03016, 2023.

Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking models for de novo molecular design. Journal of chemical information and modeling, 59(3): 1096–1108, 2019.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. Advances in neural information processing systems, 33: 1877–1901, 2020.

Stephanie Chan, Adam Santoro, Andrew Lampinen, Jane Wang, Aaditya Singh, Pierre Richemond, James McClelland, and Felix Hill. Data distributional properties drive emergent in-context learning in transformers. Advances in neural information processing systems, 35: 18878–18891, 2022.

Jorg Degen, Christof Wegscheid-Gerlach, Andrea Zaliani, and Matthias Rarey. On the art of compiling and using'drug-like'chemical fragment spaces. ChemMedChem, 3(10):1503, 2008.

Joseph L Durant, Burton A Leland, Douglas R Henry, and James G Nourse. Reoptimization of mdl keys for use in drug discovery. Journal of chemical information and computer science, 42(6): 1273–1280, 2002.

Wenhao Gao, Rocío Mercado, and Connor W Coley. Amortized tree generation for bottom-up synthesis planning and synthesizable molecular design. arXiv preprint arXiv:2110.06389, 2021.

Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor Coley. Sample efficiency matters: a benchmark for practical molecular optimization. Advances in neural information processing systems, 35: 21342–21357, 2022.

Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. What can transformers learn in-context? a case study of simple function classes. Advances in neural information processing systems, 35:30583–30598, 2022.

Qian Huang, Hongyu Ren, Peng Chen, Gregor Kr manc, Daniel Zeng, Percy S Liang, and Jure Leskovec. Prodigy: Enabling in-context learning over graphs. Advances in Neural Information Processing Systems, 36:16302–16317, 2023.

EunJeong Hwang, Veronika Thost, Shib Sankar Dasgupta, and Tengfei Ma. An analysis of virtual nodes in graph neural networks for link prediction. The first learning on graphs conference, 2022.

Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. Chemical science, 10(12):3567–3572, 2019.

Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs via the system of stochastic differential equations. In International conference on machine learning, pp. 10362–10383. PMLR, 2022.

Hyeonah Kim, Minsu Kim, Sanghyeok Choi, and Jinkyoo Park. Genetic-guided gflownets for sample efficient molecular optimization. Advances in Neural Information Processing Systems, 37: 42618–42648, 2024.

Xiangzhe Kong, Wenbing Huang, Zhixing Tan, and Yang Liu. Molecule generation by principal subgraph mining and assembling. Advances in Neural Information Processing Systems, 35: 2550–2563, 2022.

Christopher Kuenneth, Arunkumar Chitteth Rajan, Huan Tran, Lihua Chen, Chiho Kim, and Rampi Ramprasad. Polymer informatics with multi-task learning. Patterns, 2(4), 2021.

Seul Lee, Karsten Kreis, Srimukh Prasad Veccham, Meng Liu, Danny Reidenbach, Yuxing Peng, Saee Paliwal, Weili Nie, and Arash Vahdat. Genmol: A drug discovery generalist with discrete diffusion. arXiv preprint arXiv:2501.06158, 2025.

Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. arXiv preprint arXiv:2412.19437, 2024a.

Gang Liu, Tong Zhao, Jiaxin Xu, Tengfei Luo, and Meng Jiang. Graph rationalization with environment-based augmentations. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, pp. 1069–1078, 2022.

Gang Liu, Michael Sun, Wojciech Matusik, Meng Jiang, and Jie Chen. Multimodal large language models for inverse molecular design with retrosynthetic planning. arXiv preprint arXiv:2410.04223, 2024b.

Gang Liu, Jiaxin Xu, Tengfei Luo, and Meng Jiang. Graph diffusion transformers for multi-conditional molecular generation. Advances in Neural Information Processing Systems, 37: 8065–8092, 2024c.

Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? arXiv preprint arXiv:2202.12837, 2022.

Shingo Otsuka, Isao Kuwajima, Junko Hosoya, Yibin Xu, and Masayoshi Yamazaki. Polyinfo: Polymer database for polymeric materials design. 2011 International Conference on Emerging Intelligent Data and Web Technologies, pp. 22–29. IEEE, 2011.

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

Philipp Renz, Sohvi Luukkonen, and Günter Klambauer. Diverse hits in de novo molecule design: Diversity-based comparison of goal-directed generation. Journal of Chemical Information and Modeling 64(15):5756–5761, 2024.

Lloyd M Robeson. The upper bound revisited. Journal of membrane science, 320(1-2):390–400, 2008.

David Rogers and Mathew Hahn. Extended-connectivity fingerprints. Journal of chemical information and modeling, 50(5):742–754, 2010.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. Proceedings of the IEEE, 104(1): 148–175, 2015.

Aaditya K Singh, Ted Moskovitz, Sara Dragutinović, Felix Hill, Stephanie C.Y. Chan, and Andrew M Saxe. Strategy coopetition explains the emergence and transience of in-context learning. In Forty-second International Conference on Machine Learning, 2025. URL https://openreview.net/forum?id=esBoQFmD7v .

Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. Neurocomputing, 568:127063, 2024.

Michael Sun, Minghao Guo, Weize Yuan, Veronika Thost, Crystal Elaine Owens, Aristotle Franklin Grosz, Sharvaa Selvan, Katelyn Zhou, Hassan Mohiuddin, Benjamin J Pedretti, et al. Representing molecules as random walks over interpretable grammars. arXiv preprint arXiv:2403.08147, 2024.

Michael Sun, Weize Yuan, Gang Liu, Wojciech Matusik, and Jie Chen. Foundation molecular grammar: Multi-modal foundation models induce interpretable molecular graph languages. arXiv preprint arXiv:2505.22948, 2025.

A Thornton, L Robeson, B Freeman, and D Uhlmann. Polymer gas separation membrane database, 2012. URL https://research.csiro.au/virtualscreening/membrane-database-polymer-gas-separation-membranes/ .

Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. arXiv preprint arXiv:2209.14734, 2022.

Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context learning as implicit bayesian inference. arXiv preprint arXiv:2111.02080, 2021.

An Yang, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoyan Huang, Jiandong Jiang, Jianhong Tu, Jianwei Zhang, Jingren Zhou, et al. Qwen2. 5-1m technical report. arXiv preprint arXiv:2501.15383, 2025.

Botao Yu, Frazier N Baker, Ziqi Chen, Xia Ning, and Huan Sun. Llasmol: Advancing large language models for chemistry with a large-scale, comprehensive, high-quality instruction tuning dataset. arXiv preprint arXiv:2402.09391, 2024.

Barbara Zdrazil, Eloy Felix, Fiona Hunter, Emma J Manners, James Blackshaw, Sybilla Corbett, Marleen De Veij, Harris Ioannidis, David Mendez Lopez, Juan F Mosquera, et al. The chembl database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. Nucleic acids research, 52(D1):D1180–D1192, 2024.

A Discussion on Node Pair Encodings and Related Work on Fingerprints and Virtual Nodes

The node pair encodings (NPE) differ from ngerprints (Rogers & Hahn, 2010) and virtual nodes (Hwang et al., 2022) in three aspects: (1) representation and invertibility, (2) computation process, and (3) use cases. Speci cally: (1) NPE builds a motif vocabulary from data and encodes a molecule as a graph of motifs with typed edges and attachment rules, allowing the atom-level structure to be reconstructed. In contrast, Morgan ngerprints hash local atom neighborhoods into a xed-length bit vector, which is lossy and not invertible. Virtual nodes are additional nodes inserted into graphs to provide shortcuts for information exchange between nodes rather than forming a new molecular representation. If we use only virtual nodes to re-represent the molecule, the structure is not invertible. (2) NPE identi es frequent substructures from a dataset to build a discrete motif vocabulary. In contrast, Morgan ngerprints are computed directly from atom invariants with a chosen radius and bit size. Virtual nodes can be added randomly or through generic graph partitioning algorithms, which do not re ect how frequently a partition or cluster appears across the molecular dataset. (3) NPE produces a compact motif-level graph suited for generative tasks. In contrast, Morgan ngerprints produce a vector suited for predictive tasks. Virtual nodes are used to create shortcuts for sharing information between graph nodes.

B Details on DemoDiff

B.1 ICL with DemoDiff

We construct the pretraining dataset as $D_{\text{pretrain}} = \{(C_i, Q_i, X_i)\}_{i=1}^{N_{\text{pretrain}}}$. Each task contains a context of molecule–score pairs C_i, a query score Q_i, and a target molecule X_i. For each task, we rst sample a latent concept (e.g., the thermal conductivity of polymeric materials from Figure 1) from a prior $p(\theta)$, followed by a set of molecule–property pairs associated with that concept. This induces the pretraining distribution:

$$p(C, Q, X) = \int p(\theta)\, p(C \mid \theta)\, p(Q \mid C, \theta) \left[\int_{x^{1:T}} p(x^T) \prod_{t=1}^{T} p\left(x^{t-1} \mid x^t, Q, C\right) dx^{1:T} \right] d\theta, \quad (4)$$

where x^0 represents the feature of the target token. With sufficiently large and diverse pretraining data and scalable diffusion transformers (Peebles & Xie, 2023), DemoDiff can a foundation model for ICL. For a downstream task, we use (C, Q) as the prompt, and $p(\theta \mid C, Q)$ form the prompt concept. Then DemoDiff marginalizes the concept over the diffusion trajectories $x^{1:T}$:

$$p(X \mid C, Q) = \int_{\Theta} \int_{x^{1:T}} p(x^T) \prod_{t=1}^{T} p\left(x^{t-1} \mid x^t, C, Q\right) p(\theta \mid C, Q)\, dx^{1:T}\, d\theta. \quad (5)$$

B.2 Graph-level Tokens and Pretraining Loss

Graph DiTs de ne a graph-level token that concatenates the node feature with all related edge features. In molecular generation, we use a special type of edge, the null edge, to represent that there is no edge between two nodes. Thus, The feature vector (or x^0) of a graph-level token $x = \{M, \{e_j\}_{j=1}^{d}\}$ consists of three components: F_{motif} motif types, F_{bond} bond types, and F_{attach} attachment speci cations. Here F_{motif} is the size of the motif vocabulary $\mathcal{M}$, $F_{\text{bond}} = 4$ represents null, single, double, and triple bonds, and $F_{\text{attach}} = \arg\max_{M \in \mathcal{M}} |M|$ is the maximum number of

atoms in M. Eq. (3) can be decomposed as $L_{\text{pretrain}} = L_{\text{motif}} + L_{\text{bond}} + L_{\text{attach}}$. Specifically,

$$L_{\text{pretrain}} = \mathbb{E}_{q(x)} \mathbb{E}_{q(x^t|x)} \Big[\underbrace{-\log p^{\text{motif}}(m \mid x^t, C, Y)}_{L_{\text{motif}}} \underbrace{- \sum_{j=1}^{d} \log p^{\text{bond}}(b_j \mid x^t, C, Y)}_{L_{\text{bond}}} \underbrace{- \sum_{j=1}^{d} \log p^{\text{attach}}(a_j \mid x^t, C, Y)}_{L_{\text{attach}}} \Big]. \tag{6}$$

To align feature dimensions across tokens, we use the dense edge representation by treating all non-connections as null bonds. The resulting feature dimension is $F = F_{\text{motif}} + n \cdot F_{\text{bond}} + n \cdot F_{\text{attach}}$, where n is the maximum number of nodes in the motif-represented dataset. For optimization with Eq. (6), we include the null bond type in L_{bond} but exclude attachment specifications of null edges in L_{attach}.

B.3 Transition Matrices in Diffusion Models

We define the transition matrix Q that perturbs molecules at the motif level to pretrain DemoDiff. We model the joint distribution of nodes and edges with the transition matrix. It is constructed from four submatrices Q_V, Q_{EV}, Q_E, Q_{VE}, denoting transitions node $\to$ node, edge $\to$ node, edge $\to$ edge, and node $\to$ edge, respectively:

$$Q_G = \begin{bmatrix} Q_V & 1_N^0 \otimes Q_{VE} \\ 1_n \otimes Q_{EV} & 1_n \otimes 1_n \otimes Q_E \end{bmatrix}, \tag{7}$$

where $\otimes$ denotes the Kronecker product, 1_N, 1_n^0, and $1_n \otimes 1_n$ are the column vector, row vector, and all-ones matrix, respectively. Here n is the number of nodes. For edges, diffusion is applied to bond types only, while attachment attributes are optimized and predicted directly by the denoising model as in Eq. (6). For categorical sampling, we separate the unnormalized logits of node and edge from the model outputs, compute probabilities for each motif and bond individually before sampling.

To obtain the transition matrices, we use the prior from the pretraining data. The noisy distribution is defined as the marginal distributions of motif types m_V and bond types m_E. The transition matrices are defined as $Q_V = \alpha_t I + (1 - \alpha_t) 1 m_V^0$ and $Q_E = \alpha_t I + (1 - \alpha_t) 1 m_E^0$, where m^0 denotes the transpose and I is the identity matrix. We compute co-occurrence frequencies of motif and bond types in training graphs to obtain the marginal distributions m_{EV} and m_{VE}. Each row in m_{EV} gives the probability of co-occurring motifs for a bond type, and m_{VE} is its transpose. The transition matrices are then defined as $Q_{EV} = \alpha_t I + (1 - \alpha_t) 1 m_{EV}^0$ and $Q_{VE} = \alpha_t I + (1 - \alpha_t) 1 m_{VE}^0$, where α_t is cumulative noise coefficient in diffusion. The cosine schedule is chosen as $\alpha_t = \cos(0.5 \cdot (t/T + s)/(1 + s))^2$.

B.4 Details on Consistency Score

Given a query Y, demonstrations $C_i = \{(X_{ij}, Y_{ij})\}_{j=1}^{L}$ are divided into positive C^{pos}, medium C^{med}, and negative C^{neg} examples to guide ICL. For a generated molecule X, we use the Tanimoto similarity of fingerprints as the similarity measure. We compute the similarity between X and all molecules in each group and average them to obtain group-wise similarity scores

$$\text{sim}^{\text{pos}}, \text{sim}^{\text{med}}, \text{sim}^{\text{neg}} \in [0, 1].$$

Difference-based score. We compute margin differences between groups:

$$d_{\text{pos,med}} = \max(\text{sim}^{\text{pos}} - \text{sim}^{\text{med}}, 0), \quad d_{\text{med,neg}} = \max(\text{sim}^{\text{med}} - \text{sim}^{\text{neg}}, 0),$$
$$d_{\text{pos,neg}} = \max(\text{sim}^{\text{pos}} - \text{sim}^{\text{neg}}, 0).$$

The normalized difference-based score is

$$s_{di} = \min\left(\frac{d_{pos;med} + d_{med;neg} + d_{pos;neg}}{3}; 1\right):$$

In experiments, the consistency score can be computed ef ciently before applying Oracle functions. For example, we generate 1000 molecules and select the top 100 with the highest consistency scores. These molecules better follow the order of structural similarity across positive, medium, and negative examples. This removes poor generations that con ict with the demonstration semantics and increases con dence that selected molecules align with the query scores. Table 3 reports empirical improvements across task categories, each containing 4–7 tasks (appendix D.1).

Table 3: Improvement with the consistency score (average Top-10 harmonic scores).

Category	Without	With	Improvement
Drug Design	0.6230	0.7943	+27.5%
Drug MPO	0.4592	0.5400	+17.6%
Drug Rediscovery	0.4110	0.4407	+7.2%
Structure Constrained	0.5258	0.5598	+6.5%
Materials Design	0.6407	0.6724	+4.9%
Target-Based	0.7745	0.7803	+0.8%

C Details on Pretraining

The nal pretraining dataset contains 1,084,566 molecules (polymers) and 155,150 unique assays or properties, yielding 1,639,515 tasks. These are constructed from ChEMBL (Zdrazil et al., 2024) and multiple polymer data sources (Otsuka et al., 2011; Thornton et al., 2012; Kuenneth et al., 2021). Each task has a query molecule–score pair with the query score xed at 1. Up to 45 molecules are grouped into positive, medium, and negative demonstrations based on their scores. The query molecule is the target, while the query score and demonstrations serve as inputs to DemoDiff during pretraining on Eq. (3). For pretraining with a xed maximum context window, we allocate half the window to positive demonstrations and one quarter each to medium and negative demonstrations, after excluding the target molecule.

C.1 Pretraining dataset

ChEMBL dataset We constructed molecular activity contexts from the ChEMBL database (version 35), which provides a large collection of bioactivity measurements across diverse assays. ChEMBL standardizes published activity types, values, and units into a uni ed variable, pChEMBL = $-\log(\text{molar IC50, XC50, EC50, AC50, } K_i, K_d, \text{ or Potency})$. This value places different measures of half-maximal response, potency, or af nity on a comparable negative logarithmic scale. For example, an IC50 of 1 nanomolar (1×10^{-9} M) corresponds to a pChEMBL value of 9. We extracted assay-level activity values (pChEMBL). For each assay, molecules were grouped according to their recorded activities. Within each group, we selected anchor molecules with strong activity (pChEMBL > 6) as targets for building demonstrations. Each anchor was compared against all other molecules in the same assay to compute normalized distances, de ned as the relative difference between the anchor's pChEMBL value and that of the candidate molecule, converted to the range $[0, 1]$. Speci cally, for an anchor with value v_a and a candidate with value v_c, the normalized distance was given by $d = (v_1 - v_c)/10$. Based on this distance, we partitioned candidate molecules into three categories relative to the anchor. Molecules with distances in $[0, 0.25)$ correspond to candidates with activity between 75% and 100% of the anchor and were assigned to the positive context. Molecules with distances in $[0.25, 0.5)$ correspond to candidates with activity between 50% and 75% of the anchor and were assigned to the medium context. Molecules with distances in $[0.5, 1.0]$ correspond to candidates with activity below 50% of the anchor and were assigned to the negative context. From each category, we sampled up to 15 molecules to balance neighborhood size. Thus, there are up to

45 demonstration molecules for each task. Not all of them are used during pretraining due to the constraint of maximum context length. This procedure produced triplets of anchor molecules and their associated positive, medium, and negative contexts.

Polymeric materials We have polymeric material datasets from different sources, including PolyInfo (Otsuka et al., 2011), MSA (Thornton et al., 2012), and from (Kuenneth et al., 2021). We considered a wide range of polymer properties spanning several categories, including thermal properties (e.g., heat capacity, glass transition temperature, melting temperature, and thermal conductivity), electronic properties (e.g., ionization energy, electron af nity, and band gap), structural properties (e.g., density, crystallinity, and radius of gyration), and transport properties (e.g., gas diffusion, solubility, and permeability coef cients). For each property, raw values were normalized to the unit interval using min–max scaling, with logarithmic transformation applied when dynamic ranges exceeded 1000 and non-negative shifts applied when necessary. Each polymer with valid property values was treated as an anchor, and pairwise distances in normalized property space were computed against all other polymers. Candidate molecules were partitioned into positive $[0, 0.25)$, medium $[0.25, 0.5)$, and negative $[0.5, 1]$ contexts, with up to 15 examples sampled per category based on smallest distances.

C.2 Tokenizer Preparation

Algorithm 1 Node Pair Encoding (NPE) with Constraints

Require: molecule list D, motif vocabulary $M = \ ;$, max size K, ring count threshold N_{ring}
Ensure: motif vocabulary M

Initialize each molecule $X \in D$ with atom-level and ring-based motifs
Count frequencies of ring-based motifs across D
Add all periodic-table elements, polymerization "*", and top N_{ring} frequent rings to M
while $|M| < K$ **do**
  (1) Merge Neighbor: Initialize empty multiset $S \ ;$
  for each molecule $X \in D$ **do**
    for each motif m from X **do**
      for each adjacent motif m' in X such that m and m' are mergeable under structural constraints (e.g., rings treated as units) **do**
        form new motif $m \quad m \ [\ m'$
        add m to multiset S with frequency count
      end for
    end for
  end for
  (2) Frequency Selection: Find most frequent motif $m \in S$
  (3) Update Graph:
  for each molecule $X \in D$ **do**
    for each pair of adjacent motifs (m, m') in X **do**
      if their merged form equals m **then**
        replace m and m' with m in X
      end if
    end for
  end for
  Add m to M if not already in it
end while
return M

We present NPE in Algorithm 1, inspired by both the classic BPE and (Kong et al., 2022). We build the tokenizer with NPE on the pretraining data. To choose the motif vocabulary size, we analyze the number of nodes in motif-represented molecular graphs as the vocabulary size varies (Figure 8). We report mean, max, and median counts. We set $K_{ring} = K/10$, except for $K = 6000$, where $K_{ring} = 300$. When $K \quad 3000$, the mean and max node counts no longer signi cantly decrease, and the median remains unchanged. Therefore, we select $K = 3000$ with $K_{ring} = 300$ for pretraining.

Figure 8: Change in node count with varying motif vocabulary size.

Figure 9: Comparison of the number of nodes in atom- and motif-based representations for 1 million molecules from the pretraining set.

Next, Figure 9 presents the tokenization results on the pretraining dataset. log-scale distributions of motif- and atom-level node counts. Both representations exhibit heavy-tailed behavior, as shown by the rank-frequency plots and complementary cumulative distribution functions (CCDFs).

Figure 10 shows the distribution of compression ratios, de ned as $\frac{\text{Uncompressed Graph Size}}{\text{Compressed Graph Size}}$, in both linear and logarithmic scale. The ratio ranges from 1 to 40, with a median and mean around 5.5, indicating a consistent reduction in graph size by approximately a factor of ve.

Figure 11 provides a detailed analysis of the relationship between atom-level and motif-level representations. We observe a mild positive correlation: larger molecules tend to yield higher compression ratios. Notably, molecules with 150 to 200 atoms are reduced by up to a factor of 15, demonstrating ef cient compression at larger scales.

Whether NPE captures popular functional groups and reaction fragments: We evaluate the chemical validity of the motif vocabulary from the pretraining set. First, all motifs are chemically valid. Second, we compare the motifs in the vocabulary with 48 common functional groups. Forty-seven of them are present and one is missing. The groups are: