

000 001 002 003 004 005 H-MINT: MODELING POCKET-LIGAND BINDING WITH 006 HIERARCHICAL MOLECULAR INTERACTION NETWORK 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
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

ABSTRACT

Accurate molecular representations are critical for drug discovery, and a central challenge lies in capturing the chemical environment of molecular fragments, as key interactions, such as H-bond and π stacking—occur only under specific local conditions. Most existing approaches represent molecules as atom-level graphs; however, individual atoms cannot express stereochemistry, lone pairs, conjugation, and other complex features. Fragment-based methods (*e.g.*, principal subgraph or functional group libraries) fail to preserve essential information such as chirality, aromatic bond integrity, and ionic states. This work addresses these limitations from two aspects. (i) **OverlapBPE tokenization**¹. We propose a novel data-driven molecule tokenization method. Unlike existing approaches, our method allows overlapping fragments, reflecting the inherently fuzzy boundaries of small-molecule substructures and, together with enriched chemical information at the token level, thereby preserving a more complete chemical context. (ii) **h-MINT model**. We develop a hierarchical molecular interaction network capable of jointly modeling drug–target interactions at both atom and fragment levels. By supporting fragment overlaps, the model naturally accommodates the many-to-many atom–fragment mappings introduced by the OverlapBPE scheme. Extensive evaluation against state-of-the-art methods shows our method improves binding affinity prediction by 2-4% Pearson/Spearman correlation on PDBBind and LBA, enhances virtual screening by 1-3% in key metrics on DUD-E and LIT-PCBA, and achieves the best overall HTS performance on PubChem assays. Further analysis demonstrates that our method effectively captures interactive information while maintaining good generalization.

1 INTRODUCTION

Precise modeling of protein-ligand interactions is pivotal for fundamental tasks, such as binding affinity prediction and virtual screening, in early-stage drug discovery (R Laurie & Jackson, 2006). To accurately decipher these interactions, it is essential to build expressive representations to fully capture the chemical environment of the molecules. Most existing methods represent molecules as atom-level graphs (Zhou et al., 2023); however, it is questionable whether they are able to learn essential chemical information (*e.g.*, stereochemistry, lone pairs, conjugation) just from atomic tokens (Wigh et al., 2022). Another line of work employs molecular fragmentation (*e.g.*, Principal Subgraph (Kong et al., 2022b)) to preserve local contexts for atoms, in accordance with the intuition that many physicochemical properties occur at the fragment level (Murray & Rees, 2009). Nevertheless, these methods still fail to preserve crucial information like chirality, aromatic bond integrity, and ionic states, which stems from their naive partitioning of molecules into disjoint sets.

This work addresses the above limitations by integrating innovations in both molecular representations and model architectures. (i) We propose a novel and frequency-based molecular tokenization (*i.e.*, fragmentation) algorithm, **OverlapBPE**, which preserves essential chemical knowledge (*e.g.*, chirality, aromaticity, charges). Specifically, we enrich the atomic representation with their properties (*e.g.*, charges) and further incorporate 3D conformations during an iterative tokenizing process to extract frequently occurring fragments. To maintain the integrity of aromatic systems, we further enable

¹BPE (Byte Pair Encoding) (Sennrich et al., 2016) is a statistical tokenization method that iteratively merges the most frequent pairs of symbols in a corpus to generate compact and efficient subword units.

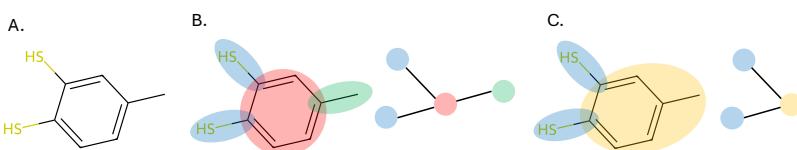


Figure 1: **Illustration of the OverlapBPE tokenization process.** (i) Starting from the molecule in A, we first extract all basic tokens from the atom graph. (ii) After identifying all basic tokens, we obtain the initial fragments (left) and token graph (right), as shown in B, which contains 4 tokens in 3 types: c1ccccc1 (freq=3778), Cc (freq=3496), and Sc (freq=637). (iii) We then enumerate all adjacent token pairs and identify the highest-frequency composite token from the final vocabulary Φ_{final} , namely Cc1ccccc1 (freq=2458). (iv) Merging c1ccccc1 and Cc into a new token, we obtain new fragments and token graph, as shown in C, containing 3 tokens in 2 types: Cc1ccccc1 (freq=2458), Sc (freq=637). (v) Continue enumerating adjacent pairs in token graph C; no matched token found in vocabulary, the algorithm terminates.

atom overlaps between mined fragments, leading to a many-to-many mapping between atoms and fragments. For example, a Naphthalene (c1c2cccc2ccc1) can be tokenized into 2 benzene rings (c1ccccc1) that share 2 aromatic C atoms. While the many-to-many mapping is necessary for aromatic integrity, it also poses an additional challenge on the model architecture, as most existing hierarchical molecular networks only support 1-1 mapping between atoms and fragments. (ii) To support our OverlapBPE tokenization, we build the hierarchical Molecular Interaction Network (**h-MINT**), which explicitly accommodates the many-to-many atom–fragment mapping. h-MINT introduces a bilevel attention mechanism that allows bidirectional information flow between atoms and overlapping fragments, and further expands fragment-level relations into atom-level geometric edges. This hierarchical yet equivariant design enables the model to capture multi-scale interaction patterns while maintaining global consistency.

Extensive experiments exhibit the superiority of our method over existing baselines, highlighting 2-4% performance gains on Pearson/Spearman correlation for binding affinity prediction, 1-3% gains on key metrics for virtual screening, and the best overall performance in HTS. Further analysis shows that our tokenization captures important inductive bias, making our model robust to noise and maintaining good generalization under different settings, indicating great potential for real-world applications.

2 BACKGROUND AND RELATED WORK

Fragment-Based Molecular Tokenization Fragmentation partitions atom-level molecular graph into coarse units that captures meaningful substructural features. Early work relied on hand-crafted junction-tree rules or predefined fragment libraries (Jin et al., 2018; 2020; Yang et al., 2021). To reduce manual bias, subsequent studies adopted unsupervised frequent-subgraph mining to construct fragmentation rules in a data-driven manner: the underlying search is NP-hard (Kuramochi & Karypis, 2001; Jazayeri & Yang, 2021), but approximate algorithms make it tractable in practice (Inokuchi et al., 2000; Yan & Han, 2002; Nijssen & Kok, 2004; Geng et al., 2023). Byte-pair encoding (BPE), which iteratively merges the most frequent token pairs to build a compact vocabulary, has also been adapted to 2D molecular data (Li & Fourches, 2021; Ucak et al., 2023; Shen & Póczos, 2024). Our work is closely related to PS-VAE (Kong et al., 2022b), which proposes a data-driven tokenization that automatically mines and merges the most frequent, maximally-sized molecular fragments (*i.e.*, *principal subgraph*). However, essential chemical information, such as stereochemistry and conjugation, is usually neglected. In contrast, our proposed OverlapBPE tackles these challenges by enriching atomic properties, involving 3D stereochemistry, and enabling overlapping between fragments.

Molecular Interaction Modeling Recent deep-learning frameworks for biomolecular modeling seek to couple fine-grained molecular representations interacting at multiple spatial resolutions. At the finest scale, atom-level graphs are usually processed with $E(n)$ -equivariant or directional message-passing networks, capturing local physics with high fidelity (Schütt et al., 2017; Satorras et al., 2021; Xu et al., 2022; Hoogeboom et al., 2022; Atz et al., 2021; Zaidi et al., 2022; Townshend et al., 2020). To reason over chemistry that spans several bonds such as aromatic conjugation,

hydrogen-bond lattices and π -stacking, researchers introduce coarser views by pooling atoms into residues or surfaces for proteins (Jin et al., 2022; Anand & Achim, 2022; Somnath et al., 2021; Wang et al., 2022), dual graphs for inverse folding (Gao et al., 2022), and functional-groups mined from small-molecule graphs (Jin et al., 2018; Kong et al., 2022b; Geng et al., 2023). Cross-molecule modules then embed ligand and receptor in a shared 3-D frame and predict poses or affinities with regression, contrastive, or diffusion objectives (Kong et al., 2022a; Luo et al., 2022; Stärk et al., 2022; Kong et al., 2024; Gao et al., 2023). Although these approaches broaden the receptive field, they often remain confined to a single resolution and are unable to enable bidirectional information flow between atoms and their corresponding substructures. Compared with prior hierarchical GNN approaches, our h-MINT introduces an atom–token overlap mechanism and expands token-level relations into atom-level geometric edges, enabling more flexible cross-scale information flow and yielding fine-grained yet globally consistent interaction modeling.

3 METHODOLOGY

3.1 OVERLAPBPE TOKENIZATION

We reveal that the failures of existing fragment-based methods stem from the naive partitioning of molecules into disjoint sets. To address these limitations, we propose a new tokenization approach that permits atom overlaps between mined fragments, enabling more expressive, coherent, and chemically meaningful molecular representations.

Atom Graph An atom graph can be represented as a property graph $G^a = (V^a, E^a)$, where $V^a = \{a_i\}$ is a set of atoms, $E^a \subset V^a \times V^a$ is the set of bonds, and each atom/bond has an associated element/bond type. A tokenization step maps a subgraph of G^a with certain atoms and bonds into a **token** $(V', E') \rightarrow f \in \Phi$, where $(V', E') = G^a[V'] \subset 2^V \times 2^E$ is an induced subgraph, and Φ is the set of tokens.

Token Graph A token graph G^f is constructed from an atom graph G^a through contracting subgraphs in G^a as a single node while preserving connectivity. To guarantee not breaking important substructures and support token overlap, we first identify a set of **basic tokens** Φ_{basic} , which includes smallest fragments that are chemically meaningful and should be preserved during partitioning. We include all single atoms, bonds, and rings collected from the training set in Φ_{basic} ² to guarantee the token set is complete, yet not all of the basic tokens will be added to the final vocabulary. In practice, we first replace all rings with tokens, then bonds, and lastly atoms to make sure all atoms and bonds are covered in the token graph G^f . In G^f , each node is a basic token, and the tokens are connected through sharing atoms (or disconnected when it is an ion).

$$G^f = (V^f, E^f), \quad \mathcal{T}^f : V^f \rightarrow \Phi_{\text{basic}} \quad (1)$$

where V^f is the set of tokens, $E^f \subseteq V^f \times V^f$ encodes their adjacency (e.g., two tokens share atoms or connected by bonds), and \mathcal{T}^f maps each fragment node to its token type in Φ_{basic} . Besides, the following tokenization process follows a bottom-up BPE merge fashion³, which also guarantees that the basic tokens will not be broken.

Frequency-Based Vocabulary Setup After we get the token graph G^f consisting of basic tokens, we iteratively discover new **composite tokens** and update the vocabulary Φ_{comp} as follows:

1. For a given token graph $G^f = (V^f, E^f)$, enumerate all adjacent fragment pairs $(f_i, f_j) \in E^f$ as composite token candidates: $\mathcal{C} = \{\text{Merge}(f_i, f_j) \mid (f_i, f_j) \in E^f\}$ where $\text{Merge}(f_i, f_j)$ denotes the operation of combining two neighboring fragments into a larger fragment. Token frequency is computed over the training corpus.
2. Select the most frequent token $f^* \in \mathcal{C}$ and add it to the vocabulary: $\Phi_{\text{comp}} \leftarrow \Phi_{\text{comp}} \cup \{f^*\}$.

²Aromatic rings are treated as single delocalized π systems to preserve conjugation.

³In BPE, tokenization starts from an initial vocabulary of basic tokens (e.g., characters or word pieces), and iteratively merges the most frequent adjacent pairs. Since merging only combines existing tokens without splitting them, the original basic tokens remain intact.

162 3. Update all fragment graphs $\{G^f\}$ in the corpus by replacing each occurrence of f^* with
 163 a new hyper node. Note that the original fragments in V^f will not be removed from G^f
 164 unless all of its adjacent composite candidates have been merged.
 165 4. Repeat steps 1–3 until a stopping criterion (e.g., iteration steps) is met.
 166 5. Filter Φ_{basic} and Φ_{comp} with minimum frequency threshold to make sure a proper vocabulary
 167 size $\Phi_{\text{final}} = \{f \in \Phi_{\text{basic}} \cup \Phi_{\text{comp}} \mid \text{freq}(f) > t\}$.
 168

169 **OverlapBPE Tokenization** After obtaining the vocabulary, we can tokenize an atom-bond graph
 170 G^a following the frequency of tokens in Φ_{final} . Figure 1 illustrates the tokenization process.
 171

172 1. Find all basic tokens Φ_{basic} from G^a and convert G^a into fragment graph G^f . For basic
 173 tokens in G^f , we preserve their token identifier if it's in Φ_{final} , otherwise, we replace it with
 174 a special identifier (e.g., <ring>) to make sure not breaking it in the following.
 175 2. For G^f , enumerate all adjacent fragment pairs $(f_i, f_j) \in E^f$ and find the matching token
 176 $f^* \in \Phi_{\text{final}}$ with the highest frequency that can be merged from pair(s) $(f_i, f_j) \in E^f$.
 177 Replace pair(s) with f^* and update G^f . Note that a token $f_i \in N^f$ will not be removed
 178 from G^f unless f_i and all its adjacent pairs $(f_i, f_j) \in E^f$ have been merged.
 179 3. Repeat step 2 until no fragment pairs match tokens in Φ_{final} .
 180

181 **Chemical Information Incorporation** Our vocabulary can be easily extended to incorporate
 182 domain knowledge. (i) To distinguish **Chirality**, the algorithm operates on 2D molecular graphs
 183 augmented with 3D conformer information. The use of 3D coordinates ensures that stereochemistry
 184 is preserved during tokenization. As a result, each token is assigned a unique isomeric SMILES string
 185 as its vocabulary identifier, encoding both atomic connectivity and chirality. For example, L-lactic
 186 and R-lactic are represented by C[C@H](O)C(=O)O and C[C@H](O)C(=O)O, respectively. (ii)
 187 **Aromatic integrity** is ensured by two complementary mechanisms. First, aromatic rings are treated
 188 as indivisible basic tokens, which are preserved throughout the bottom-up merge process. Second, the
 189 overlapping tokenization strategy allows for progressive merging of neighboring tokens, enabling the
 190 discovery of extended conjugated systems while maintaining aromatic consistency. (iii) To properly
 191 encode **atomic attributes**, such as charged and aromatic atoms, we assign explicit identifiers to
 192 atoms with formal charges and/or aromatic participation. For instance, [Cl-] denotes a negatively
 193 charged chlorine atom, while [n+] indicates a positively charged aromatic nitrogen. In contrast
 194 to standard SMILES, which often omit such details when representing isolated atoms, our token
 195 identifiers explicitly preserve these chemically significant properties.
 196

3.2 H-MINT

197 In this section, we introduce hierarchical Molecular Interaction NeTwork (h-MINT), an $SE(3)$ -
 198 equivariant graph Transformer that (i) adapts its self-attention, feed-forward, and normalization layers
 199 to maintain equivariance, and (ii) preserves atom–token overlap, allowing atoms to belong to multiple
 200 tokens (and vice versa) for flexible fragment representations. The overall architecture is in Figure 2.
 201

3.2.1 MODEL INPUT AND EMBEDDING

202 The model receives a pocket-ligand pair as input, including their atoms (V_p^a, V_l^a) ⁴, tokens (V_p^f, V_l^f) ,
 203 and atom-token mapping $\mathcal{T}_p \subseteq V_p^a \times V_p^f$, $\mathcal{T}_l \subseteq V_l^a \times V_l^f$. Note that: (i) each node in V^a or V^f has
 204 a type (e.g., element type N, token type [Cl-]). (ii) The mapping function \mathcal{T} has 2 parts to map an
 205 atom index to its associated token indices \mathcal{T}_{a2f} and vice versa \mathcal{T}_{f2a} , yet do not operate on atom and
 206 token types. Thus, the embeddings can be obtained:
 207

$$H_p^0 = \text{Embed}(V_p^a, V_p^f, \mathcal{T}_p), \quad H_l^0 = \text{Embed}(V_l^a, V_l^f, \mathcal{T}_l), \quad (2)$$

$$H^0 = \text{Embed}(V^a) + \text{ScatterMean}(\text{Embed}(V^f), \mathcal{T}_{f2a}) + \text{Embed}(\text{Pos}(V^a)), \quad (3)$$

208 where we use `scatter_mean` to aggregate an atom's associated token types information, and
 209 $\text{Pos}(\cdot)$ simply maps an atom to its position code⁵. To collect global information from atoms and
 210

211 ⁴We use subscript p and l for pocket and ligand, superscript a and f to distinguish atom and token (fragment).

212 ⁵The position code for a residue atom is its atom name, e.g., CA → A, for a molecular atom is a special token
 213 <sm>. We leave the complete position code table in the Appendix A.1.

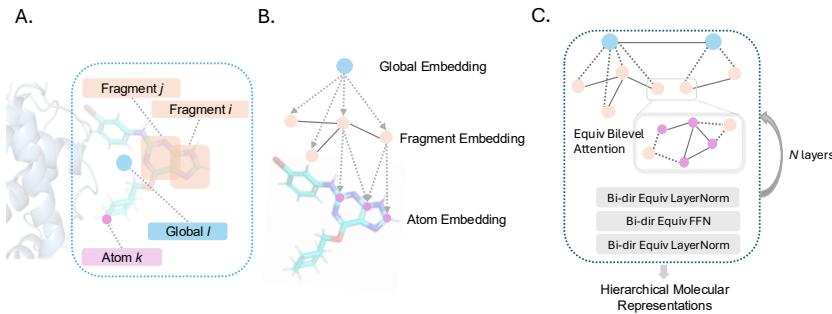
216
217
218
219
220
221
222
223
224
225
226

Figure 2: **Overall model architecture.** (A) Global node, fragment tokens, and atoms in the ligand molecule of an input pair. (B) The aggregation of fragment and global embeddings. (C) An encoder layer of h-MINT. *Note:* Solid lines indicate connection within the same level. Dashed lines indicate connections across different levels.

tokens, we also augment special `<global>` tokens to all node lists $V_p^a, V_l^a, V_p^f, V_l^f$. Some example inputs can be found in the Appendix A.1.

3.2.2 BUILDING HIERARCHICAL GRAPH

The embedding layer only encodes type and position information, and also uses $\mathcal{T}_p, \mathcal{T}_l$ to map token type embedding into atom embedding dimensions. In this section, we further consider 3D structures and build hierarchical graphs for message passing, as shown in Figure 2 B. Before we start building the graph, we need to specify that the edges are directed, which means for any token-level edge (f_i, f_j) or atom-level edge (a_i, a_j) , the former node will receive a message from the latter.

KNN Token Graph In the token level, every pair of input contains a sequence of pocket tokens and a sequence of ligand tokens, plus 2 augmented global tokens, as follows:

$$[V_p^f; V_l^f] = [f_{p,g}, f_{p,0}, f_{p,1}, f_{p,2}, \dots, f_{l,g}, f_{l,0}, f_{l,1}, f_{l,2}, \dots], \quad (4)$$

where we use f_p and f_l to represent pocket and ligand tokens, and subscription $f_{g,i}$ and $f_{i,i}$ to denote corresponding global token and i -th token. In this step, we construct a KNN graph for non-global tokens, where the distance between token f_i and f_j is the minimum distance between all their atoms,

$$\text{dist}(f_i, f_j) = \min_{a_s \in f_i, a_t \in f_j} \text{dist}(a_s, a_t), \quad (5)$$

For the two global tokens, we connect them with their following tokens to aggregate information within each pocket and ligand, and we connect the 2 global tokens to exchange pair information. Therefore we get the set of edges $E_{\text{KNN}}^f = \{(f_i, f_j) | f_i \in V_p^f \cup V_l^f, f_j \in \text{KNN}(f_i)\} \cup \{(f_g, f_i) | f_i \in V^f\} \cup \{(f_{p,g}, f_{l,g})\}$. For simplicity, we use $\text{KNN}(f_i) = \{f_j | (f_i, f_j) \in E_{\text{KNN}}^f\}$ to denote neighborhood of f_i .

Token-expanded Atom Graph The atom graph is constructed through expanding each token-level edge from E_{KNN}^f into several atom-level edges. More specifically, for a token-level edge in $(f_i, f_j) \in E_{\text{KNN}}^f$, we connect each atom with atoms from the other token. In practice, to control the number of atom-level edges, for an atom a_i , we connect $a_s \in f_i$ with the k-nearest atoms in f_j . Through this expansion, we obtain a set of atom-level edges that contain short-range interactions within atoms' neighborhoods, as well as long-range interactions bridged by token-level edges. Finally, we get the set of atom-level edges $E_{\text{knn}}^a = \{(a_s, a_t) | a_s \in f_i, a_t \in \text{knn}(f_j, a_s), (f_i, f_j) \in E_{\text{KNN}}^f\}$, and we use $\text{knn}(f_i, f_j)$ to denote the set of atom-level edges expanded from (f_i, f_j) .

3.2.3 BILEVEL MESSAGE PASSING

Prepared with aforementioned atom embeddings $[H_p^0; H_l^0]$ and token-grouped atom-level edges E_{knn}^a , we now introduce our bilevel message passing through equivariant bilevel graph attention.

270 **Equivariant Bilevel Graph Attention** For a token edge $(f_i, f_j) \in E_{\text{KNN}}^f$ and their expanded atom
 271 edges $\{(a_s, a_t) | a_s \in f_i, a_t \in \text{knn}(f_j, a_s)\}$, we first compute *atom-level cross attention* based on
 272 input atom embeddings H^{l-1} , where l is the current layer:
 273

$$[\mathbf{Q}^l; \mathbf{K}^l; \mathbf{V}^l] = \text{Linear}(H^{l-1}), \quad (6)$$

$$\text{Score} : S_{i,j}[a_s, a_t] = \text{MLP}(\mathbf{Q}^l[a_s], \mathbf{K}^l[a_t], \text{RBF}(D[a_s, a_t]), \mathbf{e}_{i,j}), \quad (7)$$

$$\text{Attention Weight} : \alpha_{i,j}[a_s, a_t] = \text{Softmax}_{a_t \in \text{knn}(f_j, a_s)}(S_{i,j}[a_s, a_t]W_\alpha), \quad (8)$$

278 where $\mathbf{Q}^l[a_s], \mathbf{K}^l[a_t]$ are query and key vectors of a_s, a_t , $\text{RBF}(D[a_s, a_t])$ embeds the relative position
 279 of a_t to a_s , $\mathbf{e}_{i,j}$ is the edge type embedding of (f_i, f_j) ⁶, and W_α is to project the scores into scalars.
 280 Through Softmax in Eq. 8, a_s is able to aggregate messages from $a_t \in \text{knn}(f_j, a_s)$, namely .
 281

282 The *token-level cross attention* is defined through aggregating all atom-level edges expanded from
 283 the token-level edge.

$$\text{Score} : S_{i,j} = \frac{1}{|\text{knn}(f_i, f_j)|} \sum_{(a_s, a_t) \in \text{knn}(f_i, f_j)} M_{i,j}[a_s, a_t], \quad (9)$$

$$\text{Attention Weight} : \beta_{i,j} = \text{Softmax}_{f_j \in \text{KNN}(f_i)}(S_{i,j}W_\beta), \quad (10)$$

284 where W_β projects the scores into scalars. Basically, $S_{i,j}$ aggregates all atom-level edges expanded
 285 from (f_i, f_j) . Then we have the following message passing and embedding update:
 286

$$\mathbf{m}_{i,j}[a_s] = \sum_{a_t \in \text{knn}(f_j, a_s)} \alpha_{i,j}[a_s, a_t] \mathbf{V}^l[a_t], \quad (11)$$

$$\mathbf{m}_i[a_s] = \sum_{f_j \in \text{KNN}(f_i)} \beta_{i,j} \text{MLP}(\mathbf{m}_{i,j}[a_s]), \quad (12)$$

$$H^l[a_s] \leftarrow H^{l-1}[a_s] + \text{ScatterMean}(\mathbf{m}_i[a_s], \mathcal{T}_{f2a}). \quad (13)$$

297 where $\mathbf{V}^l[a_t]$ is the value vector of a_t at layer- l . Stacking equivariant bilevel attention, equivariant
 298 feed-forward layer, and equivariant layer normalization together, we obtain an equivariant graph
 299 Transformer layer, as shown in Figure 2 C. For more details, please refer to Appendix A.2.

4 EXPERIMENTS AND RESULTS

300 In this section, we evaluate our OverlapBPE and h-MINT model in two fundamental drug discovery
 301 tasks: binding affinity prediction (Section 4.1) and virtual screening (Section 4.2). We conduct further
 302 experiments and case study for OverlapBPE in incorporating chemical information and representing
 303 chirality in Section 4.3.

4.1 BINDING AFFINITY PREDICTION

304 **Task Definition** Given the 3D structure of a pocket-ligand pair, the task is to predict the binding
 305 affinity, *i.e.*, change in free energy upon binding. The input is a complex (p, l) with 3D structure, and
 306 the output is a real value y . Performance is assessed with regression metrics such as root-mean-square
 307 error (RMSE) and Pearson/Spearman correlation against experimental affinities.

308 **Setup** We follow Somnath et al. (2021); Wang et al. to conduct experiments on the well-established
 309 **PDBBind** benchmark (Wang et al., 2005) and split the 4,709 complexes according to sequence identity
 310 of the protein using a 30% threshold. We also employ the **LBA** dataset with its predefined splits
 311 from the Atom3D benchmark (Townshend et al., 2020), comprising 3,507 protein-ligand complexes
 312 for training, 466 for validation, and 490 for testing. For the baseline models, we compare against a
 313 variety of approaches (Öztürk et al., 2018; Bepler & Berger, 2019; Rao et al., 2019; Elnaggar et al.,
 314 2022; Gainza et al., 2020; Hermosilla et al., 2020; Somnath et al., 2021; Wang et al., 2022; Jing et al.,
 315 2021; Townshend et al., 2020; Schütt et al., 2017; Gasteiger et al., 2021; Zhou et al., 2023; Gao et al.;
 316 Feng et al.). Details are provided in Appendix B.1. Among these models, we primarily focus our
 317 comparison on GET (Kong et al., 2024) and its variants, as it currently achieves the best performance
 318

319 ⁶We use different edge types to distinguish intra- and inter-molecule edges.

324 Table 1: **Mean and standard deviation of three runs on the PDBBind benchmark.** The best
 325 results are marked in **bold**, and the second best are underlined. Baseline results are borrowed from
 326 (Wang et al.; Kong et al., 2024). Details in Appendix B.

	Model	RMSE \downarrow	Pearson \uparrow	Spearman \uparrow
327 328 329 330 331 332 333 334 335 336 337 338	DeepDTA	1.866 ± 0.080	0.472 ± 0.022	0.471 ± 0.024
	Bepler and Berger's	1.985 ± 0.006	0.165 ± 0.006	0.152 ± 0.024
	TAPE	1.890 ± 0.035	0.338 ± 0.044	0.286 ± 0.124
	ProtTrans	1.544 ± 0.015	0.438 ± 0.053	0.434 ± 0.058
	MaSIF	1.484 ± 0.018	0.467 ± 0.020	0.455 ± 0.014
	Separate Encoder	1.554 ± 0.016	0.414 ± 0.053	0.428 ± 0.032
	IEConv	1.464 ± 0.006	0.509 ± 0.002	0.500 ± 0.005
	Holoprot-Full Surface	1.491 ± 0.004	0.491 ± 0.014	0.482 ± 0.032
	Holoprot-Superpixel	1.455 ± 0.009	0.536 ± 0.012	0.526 ± 0.012
	ProNet-Amino Acid	1.458 ± 0.003	0.546 ± 0.007	0.550 ± 0.008
	ProNet-Backbone	1.463 ± 0.001	0.551 ± 0.005	0.551 ± 0.008
	ProNet-All-Atom	1.537 ± 0.001	0.455 ± 0.013	0.433 ± 0.009
339 340 341 342 343 344 345 346	GVP	1.594 ± 0.073	-	-
	Atom3D-3DCNN	1.416 ± 0.021	0.550 ± 0.021	0.553 ± 0.009
	Atom3D-ENN	1.568 ± 0.012	0.389 ± 0.024	0.408 ± 0.021
	Atom3D-GNN	1.601 ± 0.048	0.545 ± 0.027	0.533 ± 0.033
	GET	1.430 ± 0.007	0.586 ± 0.001	0.575 ± 0.002
	GET-Murcko	1.415 ± 0.010	0.590 ± 0.002	0.578 ± 0.003
	GET-BRICS	1.410 ± 0.008	0.592 ± 0.003	0.579 ± 0.004
	GET-PS	1.387 ± 0.015	0.601 ± 0.002	0.582 ± 0.005
	Ours	1.295 ± 0.001	0.640 ± 0.002	0.625 ± 0.002

347
 348 and shares similarity with our method. Specifically, we include *GET-PS*, *GET-Murcko*, and *GET-BRICS*, which are GET models incorporating the Principal Subgraph tokenization (Kong et al., 2022b),
 349 Bemis-Murcko scaffolding (Bemis & Murcko, 1996), and BRICS tokenization (Wegscheid-Gerlach
 350 et al.). **We also include another baseline (ESM-2+fingerprint) which incorporates ESM-2 embedding**
 351 **to represents pockets and traditional fingerprints (Morgan+ERP+Avalon) to represent molecules.**

352
 353
 354 **Results** We employ OverlapBPE on the ligand molecules while utilizing residues as tokens for
 355 pockets, and conduct evaluations on PDBbind and LBA datasets separately. Table 1 and 2 report the
 356 mean and the standard deviation of the metrics across 3 runs for the PDBbind and LBA datasets. Our
 357 model demonstrates significantly better performance over baseline methods in the binding affinity
 358 prediction task. Our improvement suggests that our tokenization and modeling may better preserve
 359 interaction-relevant chemical features. Significance test and more detailed analysis are in Appendix B.

360 4.2 STRUCTURE-BASED VIRTUAL SCREENING

361
 362 **Task Definition** Given the 3D structure of a pocket p and a library of ligands $\{l_i\}$, the task is to
 363 retrieve candidate ligands that are possible to bind with p . A Virtual Screening (VS) model learns a
 364 score s_i for each pair (p, l_i) , and ranks all candidate ligands in descending order of s_i . The evaluation
 365 metrics include the area under the ROC curve (AUC), enrichment factor at a given top- k threshold
 366 (EF@ k), and BEDROC, which emphasize early-recognition of true binders.

367
 368 **Setup** Evaluation was performed on DUD-E (Mysinger et al., 2012) (102 targets, 22,886 actives)
 369 following preprocessing by Gao et al. (2023), and additionally on LIT-PCBA (Tran-Nguyen et al.,
 370 2020), **DEKOIS** (Vogel et al., 2011), and **JACS/Merck** (Wang et al., 2015; Schindler et al., 2020) as
 371 used in (Feng et al., 2025) (results in Appendix C.4). We benchmark against classical docking tools
 372 (Halgren et al., 2004; Trott et al., 2009), early ML scoring functions (Durrant & Mccammon, 2011;
 373 Ballester et al., 2010; Stepniewska-Dziubinska et al., 2017; Zheng et al., 2019; Zhang et al., 2023),
 374 and contrastive learning frameworks DrugCLIP (Gao et al., 2023) and LigUnity (Feng et al., 2025).
 375 Since DrugCLIP and LigUnity use different training sets, for a fair comparison, we train these models
 376 and our model with PDBBind (Wang et al., 2005) only, which comprises **16,744 pocket-ligand pairs**,
 377 with any overlap with the test sets removed. Following the conventions (Feng et al., 2025; Gao et al.,
 378 2023), these models are initialized with pretrained checkpoint of UniMol (Zhou et al., 2023), with
 379 loss definitions and further implementation details in Appendix C.

378
 379 **Table 2: Mean and standard deviation of three runs on LBA prediction.** The baseline results are
 380 from Kong et al. (2024). Models with * are large pretrained models, results from (Gao et al.; Feng
 381 et al.). The best results are marked in **bold** and the second best are underlined. *Note:* To save space,
 382 we only report the best baseline settings from atom-level, fragment-level, and bi-level. Details in
 383 Appendix B.

384	Best Repr. Setting	Model	RMSE ↓	Pearson ↑	Spearman ↑
385	Atom-level	SchNet	1.357 ± 0.017	0.598 ± 0.011	0.592 ± 0.015
386		EGNN	1.358 ± 0.000	0.599 ± 0.002	0.587 ± 0.004
387		LEFTNet	1.343 ± 0.004	0.610 ± 0.004	0.598 ± 0.003
388		ET	1.367 ± 0.037	0.599 ± 0.017	0.584 ± 0.025
389		UniMol*	1.434	0.565	0.540
390		BigBind*	1.340	0.632	0.620
391		ProFSA*	1.377	0.628	0.620
392	Frag-level	GemNet	1.393 ± 0.036	0.569 ± 0.027	0.553 ± 0.026
393		Equiformer	1.350 ± 0.019	0.604 ± 0.013	0.591 ± 0.012
394		DimeNet++	1.388 ± 0.010	0.582 ± 0.009	0.574 ± 0.007
395	Bi-level	MACE	1.372 ± 0.021	0.612 ± 0.010	0.592 ± 0.010
396		GET	1.331 ± 0.008	0.618 ± 0.005	0.607 ± 0.005
397		GET-PS	<u>1.312 ± 0.016</u>	<u>0.631 ± 0.011</u>	<u>0.642 ± 0.011</u>
398	Ours	1.276 ± 0.011	0.660 ± 0.001	0.661 ± 0.001	

398 **Table 3: Zero-shot Virtual Screening on DUD-E.** Baseline results are from (Gao et al., 2023; Feng
 399 et al., 2025). Models with * are trained with PDBBind data only for fair comparison. Details in
 400 Appendix C.

401	AUC (%) ↑	BEDROC (%) ↑	0.5%	EF ↑ 1%	5%
402	Glide-SP	76.70	40.70	19.39	16.18
403	Vina	71.60	-	9.13	7.32
404	NN-score	68.30	12.20	4.16	4.02
405	RFscore	65.21	12.41	4.90	4.52
406	Pafnucy	63.11	16.50	4.24	3.86
407	OnionNet	59.71	8.62	2.84	2.84
408	Planet	71.60	-	10.23	8.83
409	DrugCLIP *	81.39	45.96	34.27	29.01
410	LigUnity *	81.69	46.01	34.44	29.07
411	Ours *	84.45	47.64	35.06	29.91
412					10.76

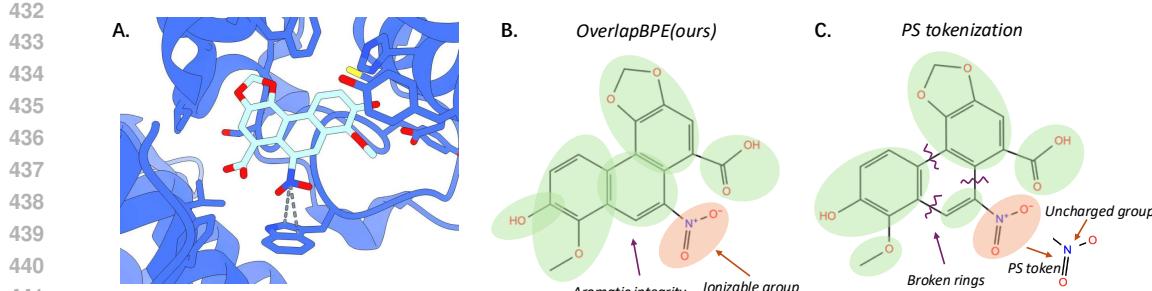
414 **Results** As shown in Table 3, our model successfully generalizes to the DUD-E benchmark under
 415 the zero-shot setting, thanks to the expressiveness of the proposed OverlapBPE and the supporting
 416 model h-MINT. Notably, our method surpasses the state-of-the-art baselines in terms of all the metrics,
 417 indicating its strong ability of capturing precise protein-ligand interaction patterns.

418 4.3 ADDITIONAL ANALYSIS OF TOKENIZATION

420 In this section, we provide additional analysis about our tokenization method’s advantages in two
 421 aspects: incorporating chemical information and representing chirality.

423 **Incorporating chemical information.** Figure 3 illustrates a case study from the LBA dataset. Our
 424 tokenization preserves the integrity of the benzene ring and retains the positive charge of $[N^+]$,
 425 which is necessary for forming the pi-cation interaction between the ligand and the protein pocket.
 426 In contrast, the PS tokenization treats $[N]$ as neutral and thus cannot capture this interaction. This
 427 difference contributes to more accurate affinity prediction, with our method achieving an error of
 428 0.56 compared to 0.67 for the PS tokenization.

429 **Representing chirality.** To further validate the effectiveness of our overlap tokenization and its capa-
 430 bility to represent chirality, we followed MolKGNN (Liu et al., 2023) and conducted high-throughput
 431 screening (HTS) experiments on PubChem assays as a binary classification task. Specifically, we
 only use the OverlapBPE tokenization with XGBoost (Chen & Guestrin, 2016) and do not adopt the



443 **Figure 3: OverlapBPE (ours) better preserves aromatic bond integrity, and ionic state.** (A)
444 An interaction formed between the ligand and the protein pocket. The ligand contains a positively
445 charged $[N^+]$, which forms two π -cation interactions with two aromatic rings in the protein pocket.
446 (B) Representation using our tokenization method. Green colors indicate fragments without charge.
447 Red colors indicate charged fragments. (C) Representation using the PS tokenization method, where
448 green and red indicate fragments as above, but the $[N^+]$ charge is not preserved, and some rings are
449 broken.

450
451
452 h-MINT model for two reasons: (i) the dataset lacks target structures, and (ii) HTS places stringent
453 demands on efficiency. We evaluate OverlapBPE on 8 PubChem HTS assays in Table 4. The baselines
454 include strong atom-level GNN models such as SchNet (Schütt et al., 2017), SphereNet (Liu et al.,
455 2022), ChiRo (Adams et al., 2021), KerGNNs (Feng et al., 2022) and MolKGNN (Liu et al., 2023).
456 For OverlapBPE, we investigate two variants: a *chiral* vocabulary that preserves chirality in tokens,
457 and a *non-chiral* vocabulary that omits stereochemical information. All baselines rely on extensive
458 molecular features and complex model architectures, whereas our method only utilizes tokenized
459 bag-of-words features combined with XGBoost for classification. We report early-enrichment perfor-
460 mance using $\text{logAUC}_{[0.001, 0.1]}$ in Table 4 and can observe that: (i) Our chiral method significantly
461 outperforms non-chiral one, showing the importance of chirality information in HTS tasks; (ii) Our
462 method outperforms all baselines in average ranking, and even exceeds ChiRo and MolKGNN, which
463 are designed to represent molecular chirality; (iii) Leveraging XGBoost’s lightweight feature, our
464 method completes training and prediction within 1 second.

465 **Table 4: Early-enrichment performance on PubChem HTS assays.** Metric is $\text{logAUC}_{[0.001, 0.1]}$
466 (higher is better). **Ours (chiral)** preserves the chirality of tokens and **Ours (non-chiral)** does not.
467 Bold numbers indicate the best method per dataset. Baseline results are taken from Liu et al. (2023).

468

PubChem AID	MolKGNN	SchNet	SphereNet	DimeNet++	ChiRo	KerGNN	Ours (chiral)	Ours (non-chiral)
435008	0.255	0.187	0.215	0.203	0.168	0.147	0.221	0.211
1798	0.174	0.195	0.196	0.208	0.165	0.078	0.217	0.282
435034	0.227	0.246	0.230	0.235	0.211	0.179	0.281	0.261
2258	0.301	0.240	0.380	0.340	0.251	0.195	0.265	0.246
463087	0.390	0.332	0.399	0.389	0.258	0.150	0.338	0.322
488997	0.303	0.319	0.309	0.315	0.193	0.081	0.384	0.376
2689	0.415	0.324	0.401	0.367	0.351	0.264	0.348	0.343
485290	0.498	0.333	0.450	0.463	0.295	0.223	0.474	0.341
Avg. Rank	3.250	5.250	3.125	3.375	6.375	8.000	2.625	4.000

5 DISCUSSION AND FUTURE WORK

482 We introduced OverlapBPE and h-MINT, an efficient tokenization-plus-learning framework for
483 protein–ligand interactions that achieves superior performance on virtual screening and affinity
484 prediction. For future work, we would like to explore ways to incorporate other domain knowledge
485 (e.g., PAINS) that is infrequent but crucial for drug discovery. Besides, we plan to extend h-MINT to
486 docking pose prediction and broader structure-based drug design.

486 ETHICS STATEMENT
487488 Small-molecule modeling plays a critical role in drug discovery, with broad potential applications in
489 therapeutic development, virtual screening, and rational design of ligands targeting protein pockets.
490 Advances in representation learning and interaction modeling offer new opportunities to accelerate
491 discovery and improve our understanding of molecular interactions, which may positively impact
492 medicine, biotechnology, and related fields.493 At the same time, we recognize that such computational approaches also carry potential risks,
494 particularly regarding misuse in unsafe or unethical drug design. To mitigate these risks, this study is
495 conducted exclusively on publicly available datasets and strictly follows established ethical guidelines.
496 We advocate for the responsible research and application of molecular modeling methods to ensure
497 their development contributes to societal benefit.499 REPRODUCIBILITY STATEMENT
500501 We ensure that the training data, training and inference procedures, and result evaluations are all
502 reproducible. The appendix provides all necessary details and offers a comprehensive explanation of
503 each component of this work. The datasets used are publicly available, and the model implementation
504 is based on the open-source GET (Kong et al., 2024), LigUnity (Feng et al., 2025) and MolKGNN
505 (Liu et al., 2023) codebases. The code and models used for evaluation are also publicly accessible
506 and cited in the appendix. Furthermore, we describe the training hyperparameters in detail in the
507 appendix, thereby ensuring that the entire experimental process is fully reproducible.509 REFERENCES
510

511 Keir Adams, Lagnajit Pattanaik, and Connor W Coley. Learning 3d representations of molecular
512 chirality with invariance to bond rotations. *arXiv preprint arXiv:2110.04383*, 2021.

513

514 Namrata Anand and Tudor Achim. Protein structure and sequence generation with equivariant
515 denoising diffusion probabilistic models. *arXiv preprint arXiv:2205.15019*, 2022.

516

517 Kenneth Atz, Francesca Grisoni, and Gisbert Schneider. Geometric deep learning on molecular
518 representations. *Nature Machine Intelligence*, 3(12):1023–1032, 2021.

519

520 Ballester, P., J., Mitchell, J., and B. A machine learning approach to predicting protein-ligand binding
521 affinity with applications to molecular docking. *Bioinformatics Oxford*, 2010.

522

523 Ilyes Batatia, David P Kovacs, Gregor Simm, Christoph Ortner, and Gábor Csányi. Mace: Higher
524 order equivariant message passing neural networks for fast and accurate force fields. *Advances in
525 Neural Information Processing Systems*, 35:11423–11436, 2022.

526

527 Guy W Bemis and Mark A Murcko. The properties of known drugs. 1. molecular frameworks.
528 *Journal of medicinal chemistry*, 39(15):2887–2893, 1996.

529

530 Tristan Bepler and Bonnie Berger. Learning protein sequence embeddings using information from
531 structure. *arXiv preprint arXiv:1902.08661*, 2019.

532

533 Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the
534 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794,
535 2016.

536

537 Weitao Du, Yuanqi Du, Limei Wang, Dieqiao Feng, Guifeng Wang, Shuiwang Ji, Carla Gomes, and
538 Zhi-Ming Ma. A new perspective on building efficient and expressive 3d equivariant graph neural
539 networks. *arXiv preprint arXiv:2304.04757*, 2023.

J. D. Durrant and J. A. Mccammon. Nnscore 2.0: A neural-network receptor–ligand scoring function.
Journal of Chemical Information & Modeling, 51(11):2897, 2011.

540 Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom
 541 Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, Debsindhu Bhowmik, and Burkhard
 542 Rost. Prottrans: Toward understanding the language of life through self-supervised learning.
 543 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):7112–7127, 2022. doi:
 544 10.1109/TPAMI.2021.3095381.

545 Aosong Feng, Chenyu You, Shiqiang Wang, and Leandros Tassiulas. Kergnns: Interpretable graph
 546 neural networks with graph kernels. In *Proceedings of the AAAI conference on artificial intelligence*,
 547 volume 36, pp. 6614–6622, 2022.

548 Bin Feng, Zijing Liu, Mingjun Yang, Junjie Zou, He Cao, Yu Li, Lei Zhang, and Sheng Wang. A
 549 foundation model for protein-ligand affinity prediction through jointly optimizing virtual screening
 550 and hit-to-lead optimization. *bioRxiv*, pp. 2025–02, 2025.

552 Shikun Feng, Minghao Li, Yinjun Jia, Wei-Ying Ma, and Yanyan Lan. Protein-ligand binding
 553 representation learning from fine-grained interactions. In *The Twelfth International Conference on
 554 Learning Representations*.

555 Pablo Gainza, Freyr Sverrisson, Frederico Monti, Emanuele Rodola, D Boscaini, MM Bronstein, and
 556 BE Correia. Deciphering interaction fingerprints from protein molecular surfaces using geometric
 557 deep learning. *Nature Methods*, 17(2):184–192, 2020.

559 Bowen Gao, Yinjun Jia, YuanLe Mo, Yuyan Ni, Wei-Ying Ma, Zhi-Ming Ma, and Yanyan Lan.
 560 Self-supervised pocket pretraining via protein fragment-surroundings alignment. In *The Twelfth
 561 International Conference on Learning Representations*.

562 Bowen Gao, Bo Qiang, Haichuan Tan, Yinjun Jia, Minsi Ren, Minsi Lu, Jingjing Liu, Wei-Ying
 563 Ma, and Yanyan Lan. Drugclip: Contrastive protein-molecule representation learning for virtual
 564 screening. *Advances in Neural Information Processing Systems*, 36:44595–44614, 2023.

566 Zhangyang Gao, Cheng Tan, Pablo Chacón, and Stan Z Li. Pifold: Toward effective and efficient
 567 protein inverse folding. *arXiv preprint arXiv:2209.12643*, 2022.

568 Johannes Gasteiger, Shankari Giri, Johannes T Margraf, and Stephan Günnemann. Fast and
 569 uncertainty-aware directional message passing for non-equilibrium molecules. *arXiv preprint
 570 arXiv:2011.14115*, 2020.

571 Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional
 572 graph neural networks for molecules. *Advances in Neural Information Processing Systems*, 34:
 573 6790–6802, 2021.

575 Zijie Geng, Shufang Xie, Yingce Xia, Lijun Wu, Tao Qin, Jie Wang, Yongdong Zhang, Feng Wu, and
 576 Tie-Yan Liu. De novo molecular generation via connection-aware motif mining. *arXiv preprint
 577 arXiv:2302.01129*, 2023.

578 Thomas A. Halgren, Robert B. Murphy, Richard A. Friesner, Hege S. Beard, Leah L. Frye, W. Thomas
 579 Pollard, and Jay L. Banks. Glide: a new approach for rapid, accurate docking and scoring. 2.
 580 enrichment factors in database screening. *Journal of Medicinal Chemistry*, 47(7):1750–1759,
 581 2004.

583 Pedro Hermosilla, Marco Schäfer, Matěj Lang, Gloria Fackelmann, Pere Pau Vázquez, Barbora
 584 Kozlíková, Michael Krone, Tobias Ritschel, and Timo Ropinski. Intrinsic-extrinsic convolution
 585 and pooling for learning on 3d protein structures. *arXiv preprint arXiv:2007.06252*, 2020.

586 Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion
 587 for molecule generation in 3d. In *International Conference on Machine Learning*, pp. 8867–8887.
 588 PMLR, 2022.

589 Akihiro Inokuchi, Takashi Washio, and Hiroshi Motoda. An apriori-based algorithm for mining
 590 frequent substructures from graph data. In *European conference on principles of data mining and
 591 knowledge discovery*, pp. 13–23. Springer, 2000.

593 Ali Jazayeri and Chris Yang. Frequent subgraph mining algorithms in static and temporal graph-
 transaction settings: A survey. *IEEE Transactions on Big Data*, 2021.

594 Yinjun Jia, Bowen Gao, Jiaxin Tan, Jiqing Zheng, Xin Hong, Wenyu Zhu, Haichuan Tan, Yuan Xiao,
 595 Liping Tan, Hongyi Cai, et al. Deep contrastive learning enables genome-wide virtual screening.
 596 *bioRxiv*, pp. 2024–09, 2024.

597 Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for
 598 molecular graph generation. In *International Conference on Machine Learning*, pp. 2323–2332,
 599 2018.

600 Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Hierarchical generation of molecular graphs
 601 using structural motifs. In *International Conference on Machine Learning*, pp. 4839–4848. PMLR,
 602 2020.

603 Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Antibody-antigen docking and design via
 604 hierarchical equivariant refinement. *arXiv preprint arXiv:2207.06616*, 2022.

605 Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael JL Townshend, and Ron Dror. Learning
 606 from protein structure with geometric vector perceptrons. *arXiv preprint arXiv:2009.01411*, 2020.

607 Bowen Jing, Stephan Eismann, Pratham N Soni, and Ron O Dror. Equivariant graph neural networks
 608 for 3d macromolecular structure. *arXiv preprint arXiv:2106.03843*, 2021.

609 Xiangzhe Kong, Wenbing Huang, and Yang Liu. Conditional antibody design as 3d equivariant graph
 610 translation. *arXiv preprint arXiv:2208.06073*, 2022a.

611 Xiangzhe Kong, Wenbing Huang, Zhixing Tan, and Yang Liu. Molecule generation by principal
 612 subgraph mining and assembling. *Advances in Neural Information Processing Systems*, 35:
 613 2550–2563, 2022b.

614 Xiangzhe Kong, Wenbing Huang, and Yang Liu. Generalist equivariant transformer towards 3d
 615 molecular interaction learning. In *International Conference on Machine Learning*, pp. 25149–
 616 25175. PMLR, 2024.

617 Michihiro Kuramochi and George Karypis. Frequent subgraph discovery. In *Proceedings 2001 IEEE
 618 international conference on data mining*, pp. 313–320. IEEE, 2001.

619 Xinhao Li and Denis Fourches. Smiles pair encoding: a data-driven substructure tokenization
 620 algorithm for deep learning. *Journal of chemical information and modeling*, 61(4):1560–1569,
 621 2021.

622 Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3d atomistic
 623 graphs. *arXiv preprint arXiv:2206.11990*, 2022.

624 Yi Liu, Limei Wang, Meng Liu, Yuchao Lin, Xuan Zhang, Bora Oztekin, and Shuiwang Ji. Spherical
 625 message passing for 3d molecular graphs. In *International Conference on Learning Representations
 (ICLR)*, 2022.

626 Yunchao Lance Liu, Yu Wang, Oanh Vu, Rocco Moretti, Bobby Bodenheimer, Jens Meiler, and
 627 Tyler Derr. Interpretable chirality-aware graph neural network for quantitative structure activity
 628 relationship modeling in drug discovery. In *Proceedings of the AAAI conference on artificial
 629 intelligence*, volume 37, pp. 14356–14364, 2023.

630 Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. Antigen-specific
 631 antibody design and optimization with diffusion-based generative models. *bioRxiv*, pp. 2022–07,
 632 2022.

633 Christopher W Murray and David C Rees. The rise of fragment-based drug discovery. *Nature
 634 chemistry*, 1(3):187–192, 2009.

635 Michael M Mysinger, Michael Carchia, John J Irwin, and Brian K Shoichet. Directory of useful
 636 decoys, enhanced (dud-e): better ligands and decoys for better benchmarking. *Journal of medicinal
 637 chemistry*, 55(14):6582–6594, 2012.

638 Siegfried Nijssen and Joost N Kok. A quickstart in frequent structure mining can make a difference.
 639 In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and
 640 data mining*, pp. 647–652, 2004.

648 Hakime Öztürk, Arzucan Özgür, and Elif Ozkirimli. Deepdta: deep drug–target binding affinity
 649 prediction. *Bioinformatics*, 34(17):i821–i829, 2018.

650

651 Alasdair T R Laurie and Richard M Jackson. Methods for the prediction of protein-ligand binding
 652 sites for structure-based drug design and virtual ligand screening. *Current Protein and Peptide
 653 Science*, 7(5):395–406, 2006.

654 Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Peter Chen, John Canny, Pieter Abbeel,
 655 and Yun Song. Evaluating protein transfer learning with tape. *Advances in neural information
 656 processing systems*, 32, 2019.

657

658 Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks.
 659 In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.

660

661 Christina EM Schindler, Hannah Baumann, Andreas Blum, Dietrich Bose, Hans-Peter Buchstaller,
 662 Lars Burgdorf, Daniel Cappel, Eugene Chekler, Paul Czodrowski, Dieter Dorsch, et al. Large-scale
 663 assessment of binding free energy calculations in active drug discovery projects. *Journal of
 664 Chemical Information and Modeling*, 60(11):5457–5474, 2020.

665

666 Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Sauceda Felix, Stefan Chmiela, Alexandre
 667 Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network
 668 for modeling quantum interactions. *Advances in neural information processing systems*, 30, 2017.

669

670 Kristof T Schütt, Huziel E Sauceda, P-J Kindermans, Alexandre Tkatchenko, and K-R Müller.
 671 SchNet—A deep learning architecture for molecules and materials. *The Journal of Chemical
 672 Physics*, 148(24), 2018.

673

674 Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with
 675 subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational
 676 Linguistics (Volume 1: Long Papers)*, pp. 1715–1725, 2016.

677

678 Yuchen Shen and Barnabás Póczos. Graphbpe: Molecular graphs meet byte-pair encoding. *arXiv
 679 preprint arXiv:2407.19039*, 2024.

680

681 Vignesh Ram Somnath, Charlotte Bunne, and Andreas Krause. Multi-scale representation learning
 682 on proteins. *Advances in Neural Information Processing Systems*, 34:25244–25255, 2021.

683

684 Hannes Stärk, Dominique Beaini, Gabriele Corso, Prudencio Tossou, Christian Dallago, Stephan
 685 Günemann, and Pietro Liò. 3d infomax improves gnns for molecular property prediction. In
 686 *International Conference on Machine Learning*, pp. 20479–20502. PMLR, 2022.

687

688 Marta M. Stepniewska-Dziubinska, Piotr Zielenkiewicz, and Paweł Siedlecki. Pafnucy - a deep
 689 neural network for structure-based drug discovery. *ArXiv*, abs/1712.07042, 2017.

690

691 Philipp Thölke and Gianni De Fabritiis. Torchmd-net: equivariant transformers for neural network
 692 based molecular potentials. *arXiv preprint arXiv:2202.02541*, 2022.

693

694 Raphael JL Townshend, Martin Vögele, Patricia Suriana, Alexander Derry, Alexander Powers,
 695 Yianni Laloudakis, Sidhika Balachandar, Bowen Jing, Brandon Anderson, Stephan Eismann, et al.
 696 Atom3d: Tasks on molecules in three dimensions. *arXiv preprint arXiv:2012.04035*, 2020.

697

698 Viet-Khoa Tran-Nguyen, Célien Jacquemard, and Didier Rognan. Lit-pcba: an unbiased data set for
 699 machine learning and virtual screening. *Journal of chemical information and modeling*, 60(9):
 700 4263–4273, 2020.

701

702 Trott, Oleg, Olson, Arthur, and J. Autodock vina: Improving the speed and accuracy of docking
 703 with a new scoring function, efficient optimization, and multithreading. *J. Comput. Chem.*, 31(2):
 704 NA–NA, 2009.

705

706 Umit V Ucak, Islambek Ashyrmamatov, and Juyong Lee. Improving the quality of chemical language
 707 model outcomes with atom-in-smiles tokenization. *Journal of cheminformatics*, 15(1):55, 2023.

708

709 Simon M Vogel, Matthias R Bauer, and Frank M Boeckler. Dekois: Demanding evaluation kits
 710 for objective in silico screening a versatile tool for benchmarking docking programs and scoring
 711 functions. *Journal of chemical information and modeling*, 51(10):2650–2665, 2011.

702 Limei Wang, Haoran Liu, Yi Liu, Jerry Kurtin, and Shuiwang Ji. Learning hierarchical protein
 703 representations via complete 3d graph networks. In *The Eleventh International Conference on*
 704 *Learning Representations*.

705 Limei Wang, Haoran Liu, Yi Liu, Jerry Kurtin, and Shuiwang Ji. Learning protein representations
 706 via complete 3d graph networks. *arXiv preprint arXiv:2207.12600*, 2022.

708 Lingle Wang, Yujie Wu, Yuqing Deng, Byungchan Kim, Levi Pierce, Goran Krilov, Dmitry Lupyan,
 709 Shaughnessy Robinson, Markus K Dahlgren, Jeremy Greenwood, et al. Accurate and reliable
 710 prediction of relative ligand binding potency in prospective drug discovery by way of a modern
 711 free-energy calculation protocol and force field. *Journal of the American Chemical Society*, 137
 712 (7):2695–2703, 2015.

713 Renxiao Wang, Xueliang Fang, Yipin Lu, Chao-Yie Yang, and Shaomeng Wang. The pdbsbind
 714 database: methodologies and updates. *Journal of medicinal chemistry*, 48(12):4111–4119, 2005.

715 Christof Wegscheid-Gerlach, Andrea Zaliani, and Matthias Rarey. On the art of compiling and
 716 using ‘drug-like’ chemical fragment spaces.

718 Daniel S Wigh, Jonathan M Goodman, and Alexei A Lapkin. A review of molecular representation in
 719 the age of machine learning. *Wiley Interdisciplinary Reviews: Computational Molecular Science*,
 720 12(5):e1603, 2022.

722 Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric
 723 diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*, 2022.

724 Xifeng Yan and Jiawei Han. gspan: Graph-based substructure pattern mining. In *2002 IEEE*
 725 *International Conference on Data Mining, 2002. Proceedings.*, pp. 721–724. IEEE, 2002.

727 Soojung Yang, Doyeong Hwang, Seul Lee, Seongok Ryu, and Sung Ju Hwang. Hit and lead discovery
 728 with explorative rl and fragment-based molecule generation. *Advances in Neural Information*
 729 *Processing Systems*, 34, 2021.

730 Sheheryar Zaidi, Michael Schaarschmidt, James Martens, Hyunjik Kim, Yee Whye Teh, Alvaro
 731 Sanchez-Gonzalez, Peter Battaglia, Razvan Pascanu, and Jonathan Godwin. Pre-training via
 732 denoising for molecular property prediction. *arXiv preprint arXiv:2206.00133*, 2022.

734 Xiangying Zhang, Haotian Gao, Haojie Wang, Zhihang Chen, Zhe Zhang, Xinchong Chen, Yan
 735 Li, Yifei Qi, and Renxiao Wang. Planet: A multi-objective graph neural network model for
 736 protein–ligand binding affinity prediction, 2023. URL <https://doi.org/10.1101/2023.02.01.526585>.

738 Liangzhen Zheng, Jingrong Fan, and Yuguang Mu. Onionnet: a multiple-layer intermolecular-contact-
 739 based convolutional neural network for protein–ligand binding affinity prediction. *ACS omega*, 4
 740 (14):15956–15965, 2019.

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

746

747

748

749

750

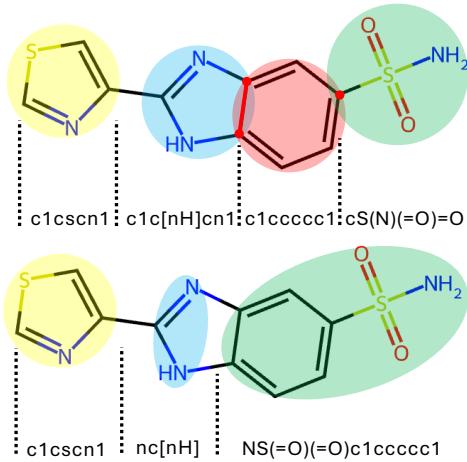
751

752

753

754

755

756 A METHOD
757758 A.1 OVERLAPBPE
759777 Figure 4: Comparison of overlap (top) and non-overlap tokenization (bottom).
778

779 **Overlap vs. Non-overlap Tokenization** In Figure 4, we show an example molecule tokenized in 2
780 different ways. We use different colors to highlight the tokenized fragments. From the top figure,
781 we can see the molecule has 4 tokens, with 3 of the tokens sharing atoms with other tokens (we also
782 highlight the shared atoms and bonds in red). From the bottom figure, we can see that an aromatic
783 bond is broken and forms 2 disjoint tokens. From this figure, we can identify the difference between
784 overlap-tokenization and non-overlap-tokenization. With overlap-tokenization better preserving the
785 local chemical environment and better respecting the fuzzy boundaries of substructures in small
786 molecules, we believe our novel OverlapBPE is potent to boost small molecular representation,
787 interaction, and even generation tasks.

788 **Example Input** Here we show some example inputs to the encoder, more specifically, the embedding
789 layer. The embedding layer takes 3 inputs: atom sequences, atom-level position codes, and
790 token sequences, for pairs of pocket-ligands.
791

$$792 [V_p^a; V_l^a] : [\langle g_atom \rangle, N, C, C, O, C, \dots, \langle g_atom \rangle, C, C, \dots] \quad (14)$$

$$793 [Pos(V_p^a); Pos(V_l^a)] : [\langle g_pos \rangle, \text{''}, A, \text{''}, \text{''}, B, \dots, \langle g_pos \rangle, sm, sm, \dots] \quad (15)$$

$$794 [V_p^f; V_l^f] : [\langle g_frag \rangle, ALA, \dots, \langle g_frag \rangle, c1cscn1, \dots] \quad (16)$$

795 **Position Code** Certain types of molecules have conventional position codes to distinguish different
796 atoms with the same element type in the same fragment, *e.g.*, CA and CB in residues. For small
797 molecules, since there are no such conventional position codes, we simply use *sm* as the position code
798 for all atoms in small molecules. In addition, we also add some special position codes for special
799 tokens, *e.g.*, *<global>*, *<mask>*, *<pad>*.

800 **Tokenization Overhead** Compared to the PS tokenizer and GET, the computational overhead
801 of OverlapBPE and h-MINT mainly stems from the repeated calculation of atoms. We performed
802 tokenization using OverlapBPE and PS on the 3,507 molecules in the LBA training set, and the
803 results are shown in Table 5.

804 For tokenization, the number of atoms processed by OverlapBPE is 1.32 times that of PS: (7.95×4.5)
805 / $(8 \times 3.39) = 1.32$. However, since tokenization can be **executed offline and is fully parallelizable**,

810
811
812 Table 5: Comparison of tokenization statistics.
813
814
815

Tokenizer	Avg # tokens / mol	Avg # atoms / token
OverlapBPE	7.95	4.5
Principal Subgraph	8	3.39

816
817 the overhead is negligible in practice. For instance, OverlapBPE only takes **7 minutes** to process
818 **47.9k molecules** using 32 CPUs for virtual screening training data.
819

820 A.2 h-MINT

821 **SE-(3) Equivariance** In this section, we provide more details about the model design. At first, h-
822 MINT follows GET’s (Kong et al., 2024) architecture with 2-channel updates: One is the equivariant
823 channel, which mainly encodes and predict the coordinates of molecules following SE-(3) symmetry.
824 The other is the invariant channel, which mainly encodes and predicts embeddings (*i.e.*, H). Thus,
825 in general, our model is an SE-(3) equivariant model. Since we only use the embedding channel in
826 this paper, we mainly show the update and message passing for the invariant channel. However, our
827 model can also be extended to SE-(3) scenarios like structure prediction and generation.
828

829 **Difference with GET** We want to emphasize the difference between h-MINT and GET. GET
830 is designed for non-overlap tokenization, while h-MINT is designed for overlap tokenization, and
831 this induces a fundamental difference in model design. For GET, atoms and tokens are in a 1-1
832 mapping, making its model design simple and straightforward. For h-MINT, atoms and tokens are
833 in a many-to-many mapping, which requires a bidirectional indexing system to convert atom-level
834 embeddings to token-level embeddings and vice versa.
835

836 In Section 3.2, we present the embedding layer, graph construction, and Bilevel Graph Attention
837 Layer. For the next, we complement other modules, including Bidirectional Equivariant Feed-Forward
838 Network (FFN) and Bidirectional Equivariant Layer Normalization (LN).
839

840 **Bidirectional Equivariant FFN** For this module, the input contains the atom-level embedding
841 H^l passed from the former module, and also the bidirectional mapping between atoms and tokens
842 $\mathcal{T}_{a2f}, \mathcal{T}_{f2a}$. The update is as follows:
843

$$H^{l,f} = \text{ScatterMean}(H^l, \mathcal{T}_{a2f}) \quad (17)$$

$$H^{l'} = \text{ScatterMean}(H^{l,f}, \mathcal{T}_{f2a}) \quad (18)$$

$$H^l \leftarrow H^l + \text{MLP}([H^l; H^{l'}; \text{RBF}(D)]), \quad (19)$$

844 where $H^{l'}$ can be regarded as token-enhanced atom representation, and D stores the pairwise atom
845 distances.
846

847 **Bidirectional Equivariant LN** This module involves normalization within each input pair, since we
848 do not use the equivariant channel. Here, we apply simple Layer Normalization to the representations.
849

$$H^l \leftarrow \frac{H^l - \mathcal{E}[H^l]}{\sqrt{\text{Var}[H^l] + \epsilon}} \cdot \sigma + \mu, \quad (20)$$

850 where σ and μ are learnable parameters, and $\mathcal{E}[\cdot]$ and $\text{Var}[\cdot]$ are used to calculate the mean and variance
851 of the variable.
852

853 B EXPERIMENTS: BINDING AFFINITY PREDICTION

856 B.1 PDBBind BENCHMARK

857 We follow the data processing of (Somnath et al., 2021; Wang et al.) to conduct experiments on
858 the PDBBind (v2019) Benchmark. More specifically, we use the split with a sequence identity of
859 30% to prevent leakage. This filtering results in 4,709 complexes, which are then split into 3,507,
860

864 466, and 490 for training, validation, and testing (Somnath et al., 2021). We directly borrow the
 865 baseline results from (Wang et al.; Kong et al., 2024). The main advantage of GET-PS and our model
 866 is tokenizing small molecules into fragments. To evaluate these models, we simply remove pairs
 867 without small molecules, which results in a slightly smaller split (4436 samples) compared to the
 868 original identity30 data split (4463 samples).

870 B.2 IMPLEMENTATION DETAILS

871 We conduct experiments on 1 RTX A6000 GPU. Each model is trained with the Adam optimizer
 872 and learning rate decay. Considering the number of tokens and atoms may vary with different input
 873 complexes, to safely and efficiently utilize the GPU memory, we implement a dynamic batch to
 874 include as many complexes as possible while not exceeding some threshold `max_n_vertex`. For graph
 875 construction, we use $k = 9$ for the KNN token graph, which means each token is connected with 9
 876 nearest tokens within a complex. And we use $k = 3$ for the token-expanded atom graph, which means
 877 for atom $a_s \in f_i$, for each KNN of the token, $f_j \in \text{KNN}(f_i)$, a_s will connect to 3 nearest atoms
 878 in f_j . We set the RBF kernel size to be 32. For baseline models, we follow the official parameter
 879 configuration. For our model, we mainly tune the learning rate ($\text{lr} \in [1e-3, 1e-4]$), final learning
 880 rate ($\text{flr} \in [1e-3, 1e-6]$), and max number of epochs ($\text{max_epoch} \in [10, 40]$). Here we list the
 881 key hyperparameters for reproducing our results on PDBBind ($\text{lr} = 1e-4$, $\text{flr} = 1e-4$, $\text{max_epoch} = 40$)
 882 and LBA ($\text{lr} = 1e-3$, $\text{flr} = 1e-4$, $\text{max_epoch} = 20$).

883 B.3 BASELINE MODELS

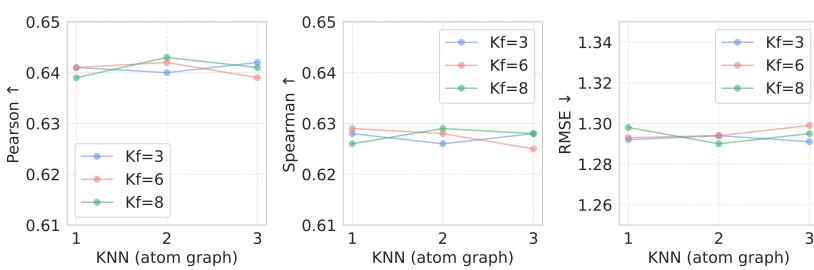
884 We include many baseline models in this section of experiments. Here we briefly introduce their
 885 methods. For the **PDBBind** benchmark, DeepDTA (Öztürk et al., 2018), Bepler and Berger's (Bepler
 886 & Berger, 2019), TAPE (Rao et al., 2019), ProtTrans (Elnaggar et al., 2022), MaSIF Gainza et al.
 887 (2020), IEConv (Hermosilla et al., 2020), Holoprot (Somnath et al., 2021), and ProNet (Wang et al.)
 888 use separate encoders for pockets and ligands. GVP (Jing et al., 2020), Atom3D (Townshend et al.,
 889 2020), and GET (Kong et al., 2024) instead use a joint encoder for pockets and ligands. Inspired by
 890 the good performance and trends in joint encoder models, we also adopt a joint encoder architecture.

891 For the **LBA** dataset, SchNet (Schütt et al., 2018), DimeNet++ (Gasteiger et al., 2020), GemNet
 892 (Gasteiger et al., 2021) are invariant models based on invariant geometric features (*i.e.*, distance
 893 and angle). EGNN (Satorras et al., 2021), TorchMD-Net (ET) (Thölke & De Fabritiis, 2022),
 894 and LEFTNet (Du et al., 2023) preserve equivariant features and are directly implemented on 3D
 895 coordinates. MACE (Batatia et al., 2022) and Equiformer (Liao & Smidt, 2022) utilize harmonic and
 896 irreducible representations to preserve high-order equivariant features. We also include atom-level
 897 pretrained models, UniMol (Zhou et al., 2023), ProFSA (Gao et al.) and BigBind (Feng et al.). In
 898 general, all these models mainly use their invariant channel for affinity prediction, similar to GET
 899 (Kong et al., 2024) and our model; thus, we can ignore how these models deal with equivariant
 900 features in these experiments. We borrow the baseline results mainly from GET (Kong et al., 2024),
 901 which provides a complete comparison of all the above models in 3 representation settings: atom-
 902 level, fragment-level, and bi-level. To save space, we include each model's best representation setting
 903 only.

904 B.4 ADDITIONAL EXPERIMENT ANALYSIS

905 **Significance of Improvements** Since the benchmarks (PDBBind and LBA) are relatively small, to
 906 guarantee fair comparison and consistent results, we report mean and std for 3 runs in Table 1 and 2.
 907 Besides, we also conduct a significance test on the prediction results of GET, GET-PS and our model.
 908 And the results show p -values < 0.005 for these models in both PDBBind and LBA tasks. These
 909 evidences conclude that our model performs significantly better than the strong baselines, GET, and
 910 GET-PS.

911 **PDBBind Results Analysis** From Table 1, we can mainly draw the following conclusions: (i)
 912 Joint encoder is generally better than separate encoders. (ii) GET outperforms other baselines by a
 913 wide margin due to its hierarchical modeling, and GET-PS performs even better for better alignment
 914 between residues and small molecule fragments. (iii) Our model consistently outperforms GET and
 915 GET-PS by a wide margin (4%-5% Pearson Correlation, and 5% Spearman Correlation). We are

928 **Figure 5: Parameter study of KNN graph construction.** Kf: KNN in fragment-level graph.
929

930 also surprised by the improvements, and we recognize that the improvements come from our new
931 tokenization, OverlapBPE, and our new model h-MINT.
932

933 As defined in Section 3.2.2, we employ K-Nearest Neighbors (KNN) to construct both fragment-level
934 and atom-level graphs, where the choice of the hyperparameter K plays a critical role in determining
935 the quality of graph connectivity and the overall model performance. To assess the impact of different
936 K values, we conducted a parameter study on the PDDBBind dataset. The results are illustrated in
937 Figure 5. From the results we can see that our model is not sensitive to the choice of K. Therefore,
938 selecting the parameter within a proper range can consistently yield satisfactory performance.
939

940 **LBA Results Analysis** Table 2 compares different models with atom-level, fragment-level, and
941 bi-level representations. We borrow the baseline results from (Kong et al., 2024), which compares
942 all baselines in all 3 representation settings. In this paper, we only include the baseline results in
943 their best representation settings for saving space. From this table, we can see: (i) In general, bi-level
944 representation is better than atom-level or fragment-level. This is reasonable because some molecular
945 interactions happen on atoms (H-bond) and some happen on fragments (π stacking). Thus, we believe
946 integrating bi-level information is crucial for modeling molecular interactions. (ii) GET and GET-PS
947 still outperform other baselines by a wide margin (2% Pearson and Spearman Correlation), which
948 demonstrates the effectiveness of their unified representations and model design. (iii) Last, our model
949 outperforms GET and GET-PS even more than their improvements (3% Pearson Correlation, 2%
950 Spearman Correlation). This result again validates the effectiveness of our OverlapBPE tokenization
951 and our new model for many-to-many mapping between atoms and tokens.
952

953 C EXPERIMENTS: VIRTUAL SCREENING

954 C.1 DETAILS ON LOSS FUNCTION

955 According to LigUnity (Feng et al., 2025), we optimise a composite objective
956

$$\mathcal{L} = \underbrace{(\mathcal{L}_{p \rightarrow l} + \mathcal{L}_{l \rightarrow p}) + \mathcal{L}_{\text{rank}}}_{\mathcal{L}_{\text{LigUnity}}} + \lambda_{\text{mse}} \mathcal{L}_{\text{mse}}, \quad (1)$$

957 which extends the original loss with an additional regression term.
958

959 **Contrastive retrieval losses.** For a mini-batch of B pocket embeddings p_i and ligand embeddings l_j
960 we define
961

$$\mathcal{L}_{p \rightarrow l} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\tau \langle p_i, l_i \rangle)}{\sum_{j=1}^B \exp(\tau \langle p_i, l_j \rangle)}, \quad \mathcal{L}_{l \rightarrow p} \text{ sym.}, \quad (2)$$

962 **Listwise ranking loss.** Given a pocket i with M_i ligands sorted by experimental affinity $\pi_1 \succ \dots \succ$
963 π_{M_i} ,

$$\mathcal{L}_{\text{rank}} = - \sum_{k=1}^{M_i-1} \mu_k \log \frac{\exp(\tau \langle p_i, l_{\pi_k} \rangle)}{\sum_{t=k}^{M_i} \exp(\tau \langle p_i, l_{\pi_t} \rangle)}, \quad (3)$$

964 with positional weights $\mu_k = \frac{1}{\log(k+1)}$.
965

972
 973 **MSE loss.** Similar to the regression loss implemented in the LigUnity paper, let \hat{a}_{ij} be the predicted
 974 activity for pair (i, j) and a_{ij} the ground truth. With \mathcal{P} the positive set, \mathcal{N} represents a 20 %
 975 subsample of negatives, pocket-wise weakest positive $a_{\min,i}$ and safety margin δ , which was set to
 976 2.0:
 977

$$\mathcal{L}_{\text{mse}} = \frac{1}{|\mathcal{P}| + |\mathcal{N}|} \left(\sum_{(i,j) \in \mathcal{P}} (\hat{a}_{ij} - a_{ij})^2 + \sum_{(i,j) \in \mathcal{N}} \left[\max(0, \hat{a}_{ij} - (a_{\min,i} - \delta)) \right]^2 \right). \quad (4)$$

980 Throughout this paper, we fix the weights to $\lambda_{\text{mse}} = 2$. Empirically, the extra MSE term accelerates
 981 convergence and mitigates the overfit on the training dataset.
 982

983 C.2 EVALUATION METRICS

984 We assess model performance using the following metrics:
 985

- 986 • **AUC-ROC (Area Under the Receiver Operating Characteristic curve):** measures the
 987 probability that a randomly chosen active compound is ranked higher than a randomly
 988 chosen inactive one. In our paper, we use AUC to denote it.
 989
- 990 • **BEDROC (Boltzmann-Enhanced Discrimination of Receiver Operating Character-
 991 istic):** emphasizes early recognition of actives by applying an exponential weighting to
 992 the ROC curve, controlled by a tunable parameter. Following previous works, we set the
 993 parameter to 80.5.
 994
- 995 • **Enrichment Factor (EF):** quantifies the fold-increase in actives found among the top
 996 percentile of the ranked library relative to random selection, reported here at 0.5%, 1%, and
 997 5%.
 998

999 C.3 IMPLEMENTATION DETAILS

1000 The implementation of virtual screening experiments includes two parts: training the baselines and
 1001 finetuning our model. All the models were trained on 1 RTX 6000 GPU with 24 GB memory. For the
 1002 baseline models, we retrained LigUnity (Feng et al., 2025) on the PDDBind dataset. To ensure a fair
 1003 evaluation, we excluded all the samples that exist in any of the test datasets. We also followed their
 1004 papers’ original parameters. To evaluate, we averaged the weights of the last 3 model checkpoints.
 1005 We finetuned our h-MINT model on the same PDDBind dataset. Namely, the training parameters
 1006 were similar to LigUnity, with learning rate = $1e-4$, warmup ratio = 0.06, and the maximum number
 1007 of ligands selected for each pocket was 16. However, we used 32-bit precision for training rather
 1008 than the original 16-bit, changed the batch size to 96, and set the gradient clip to 10. We trained the
 1009 model for 100 epochs initially, and observed that it converged at around the 25th epoch. Therefore,
 1010 we trained it for 25 epochs and averaged the last 3 checkpoints for evaluation purposes.
 1011

1012 C.4 EXPERIMENTS ON LIT-PCBA, DEKOIS AND JACS/MERCK DATASET

1013 C.4.1 LIT-PCBA DATASET

1014 The LIT-PCBA dataset (Tran-Nguyen et al., 2020) is constructed from dose–response PubChem
 1015 bioassays to mitigate the target and decoy selection biases found in other benchmarks. It comprises
 1016 15 protein targets with 7844 experimentally confirmed actives and 407381 inactive, which reflects
 1017 realistic hit rates in high-throughput screening.
 1018

1019 C.4.2 ADDITIONAL RESULTS AND SIGNIFICANCE TEST

1020 Table 6 reports AUC, BEDROC, and enrichment factors (EF) on LIT-PCBA in the zero-shot setting.
 1021 Namely, as the original LigUnity was trained on another larger dataset Feng et al. (2025), we did not
 1022 directly compare our method with it. However, we also used their provided checkpoint and evaluated
 1023 its performance on the PCBA dataset. It could be concluded that LigUnity trained on its full dataset
 1024 leads the zero-shot performance, which also highlights the value of larger training datasets, similar
 1025 to Table 3. When constrained to PDDBind alone, our approach outperforms LigUnity across most
 1026 metrics, achieving an AUC of 57.77% and $\text{EF}_{0.5}$ of 7.01. These results strongly demonstrate its
 1027

1026 Table 6: **Zero-shot Virtual Screening on LIT-PCBA.** Baseline results are from Gao et al. (2023);
 1027 Feng et al. (2025); Jia et al. (2024). Models with * are trained with PDBBind data only. Bold numbers
 1028 indicate the best performance.

	AUC (%) ↑	BEDROC (%) ↑	0.5%	EF ↑ 1%	5%
Surflex	51.47	-	-	2.50	-
Glide-SP	53.15	4.00	3.17	3.41	2.01
Planet	57.31	-	4.64	3.87	2.43
Gnina	60.93	5.40	-	4.63	-
DeepDTA	56.27	2.53	-	1.47	-
BigBind	60.80	-	-	3.82	-
DrugCLIP *	58.15	4.12	4.11	3.08	2.27
LigUnity *	57.61	4.34	4.06	3.03	2.25
Ours *	57.77	6.27	7.01	5.20	2.18

Table 7: **Zero-shot Virtual Screening on DEKOIS.**

	AUC (%) ↑	BEDROC (%) ↑	0.5%	EF ↑ 1%	5%
LigUnity	76.92	47.20	18.57	16.25	8.21
Ours	81.05	47.71	18.085	16.77	8.74

1050 robustness under limited training data and the effectiveness of our efficient tokenization strategy. In
 1051 addition, we also provides comparison with LigUnity following (Feng et al., 2025) in Table 7 and 8.
 1052 These results validate that h-MINT generalizes effectively across datasets and tasks, including both
 1053 affinity-ranking and virtual-screening benchmarks.

1054 To demonstrate the significance of our results, we conducted statistical significance tests for all
 1055 benchmarks (DUDE, PCBA, DEKOIS, and FEP). For all comparisons between h-MINT(ours) vs.
 1056 LigUnity, we obtained p -values < 0.005 , demonstrating that the improvements are statistically
 1057 significant and consistent, rather than due to random variation.

C.5 ABLATION OF MSE LOSS OVER LIGUNITY

1061 To isolate the contribution of this loss, we trained LigUnity with additional MSE loss (exact same
 1062 combined loss as ours) under identical settings on the PDBBind training set, as shown in Table 9. The
 1063 results on DUDE and LIT-PCBA are shown in the table above, which confirms the following findings:
 1064 (i). Effect of the proposed auxiliary loss. Using our proposed auxiliary loss consistently improves
 1065 LigUnity across almost all metrics on both datasets. For example, on DUDE, AUC improves from
 1066 81.69 to 82.57, and BEDROC from 46.01 to 47.58. On LIT-PCBA, early-recognition metrics show
 1067 noticeable gains as well. This confirms that additional loss is beneficial and strengthens the model’s
 1068 scoring ability. (ii). Advantage of h-MINT (LigUnity+MSE vs Ours). Our model further improves
 1069 over LigUnity+MSE on most metrics. Gains are particularly clear in early-recognition measures
 1070 such as BEDROC and EF0.01, which are widely regarded as key metrics for virtual screening. These
 1071 results confirm that: our proposed regression loss is effective, but our h-MINT architecture delivers
 1072 additional, consistent boosts beyond what the regression loss alone can offer. Thus, the comparison
 1073 with LigUnity is fair, and the observed improvements come from both components of our method.

Table 8: **Zero-shot Virtual Screening on JACS/Merck.**

	r^2
LigUnity	0.173
Ours	0.216

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
Table 9: **Ablation of MSE Loss over LigUnity**

Dataset	Model	AUC (%) \uparrow	BEDROC (%) \uparrow	0.5%	EF \uparrow 1%	5%
DUDE	LigUnity	81.69	46.01	34.44	29.07	10.26
	LigUnity+MSE	82.57	47.58	35.83	29.77	10.70
	Ours	84.47	47.65	35.06	29.90	10.76
LIT-PCBA	LigUnity	57.61	4.34	4.07	3.04	2.26
	LigUnity+MSE	57.68	5.64	6.50	4.22	2.14
	Ours	57.77	6.21	7.01	5.20	2.18

Table 10: **Vocabulary Statistics on LBA training set.**

min_freq	# tokens (basic / composite / all)	avg token size (basic / composite / all)
200	41 / 52 / 93	2.56 / 5.81 / 4.38
100	54 / 80 / 134	2.87 / 6.53 / 5.05
50	74 / 137 / 211	3.11 / 7.44 / 5.92
20	94 / 200 / 294	3.28 / 8.11 / 6.56
10	112 / 400 / 512	3.54 / 9.55 / 8.23

C.6 COMPUTATIONAL OVERHEAD

OverlapBPE duplicates certain atoms during tokenization to maintain the continuity and integrity of chemical substructures. Because of these overlapping tokens, the final number of atoms becomes roughly 1.32 times the original. Despite this, both tokenization and graph construction are highly parallelizable and can be performed fully offline in preprocessing. In practice, the overhead is negligible: OverlapBPE only takes 7 minutes to process 47.9k molecules using 32 CPUs for virtual screening training data.

During training and inference, the main computational cost comes from the underlying UniMol encoder. h-MINT functions as a light-weight adapter, and the runtime difference compared with LigUnity is minimal: our’s training time is $\times 1.12$ than LigUnity, and our’s inference time is $\times 1.07$ than LigUnity. Therefore, although h-MINT introduces a richer atom-fragment representation, the parallel and offline preprocessing ensures that the runtime during the actual virtual screening pipeline remains nearly unchanged.

D ADDITIONAL ANALYSIS AND EXPERIMENTS

D.1 VOCABULARY STATISTICS

We provide vocabulary statistics on LBA training set in Table 10. For all the datasets, we extract vocabularies solely from the training set, without introducing additional data.

D.2 REPRESENTING CHIRALITY FOR HTS

We evaluate whether explicitly encoding stereochemistry in our fragment representation improves early enrichment on eight PubChem HTS assays. In contrast to structure-based virtual screening in Section 4.2, HTS assay data do not include explicit target information. Each assay contains active and inactive compounds corresponding to the same target. We construct two OverlapBPE vocabularies: (i) *Ours-non-chiral*, which omits stereochemical markers, and (ii) *Ours-chiral*, which operates on 2D graphs augmented with 3D conformer cues (used only to disambiguate stereocenters) and assigns isomeric SMILES identifiers to tokens, thereby preserving R/S and @/@@ annotations. Each molecule is tokenized into overlapping fragments; we compute sparse bag-of-tokens features and train an XGBoost logistic classifier per assay on the official train/valid/test split. We report $\text{logAUC}_{[0.001, 0.1]}$, which emphasizes the low false-positive region relevant to early enrichment.

Table 11: Ablation Study of OverlapBPE and h-MINT on LBA.

	RMSE \downarrow	Pearson \uparrow	Spearman \uparrow
GET	1.331 ± 0.008	0.618 ± 0.005	0.607 ± 0.005
GET+PS	1.312 ± 0.016	0.631 ± 0.011	0.642 ± 0.011
h-MINT+PS	1.321 ± 0.010	0.633 ± 0.007	0.641 ± 0.008
GET+OverlapBPE	N/A	N/A	N/A
Ours (h-MINT+OverlapBPE)	1.276 ± 0.011	0.660 ± 0.001	0.661 ± 0.001

D.3 ABLATION STUDY OF OVERLAPBPE AND h-MINT

OverlapBPE tokenization and overlap-compatible hierarchical interaction model, jointly form an integrated framework to tackle the challenge of fuzzy boundaries of meaningful molecular substructures in 3D molecular interaction modeling. Meaningful comparison can only be made when they’re used together, because (i) no other network architectures are available for overlapping substructures, and (ii) when non-overlapping tokenization is used, the molecular graphs for h-MINT and GET become identical. As evidence, we provide the following ablation on LBA dataset that adopts non-overlap PS tokenizer for h-MINT in Table 11.

As can be seen in this table, GET is not compatible with OverlapBPE. OverlapBPE+hMINT consistently outperforms models with PS tokenizer, demonstrating a clear gain of OverlapBPE. When non-overlap PS tokenizer is used, GET+PS and h-MINT+PS achieve similar performance as expected. The tokenizer and h-MINT architecture are both necessary to handle overlapping fragments and preserve key chemical information (atomic integrity, ionic states, chirality).

D.4 NOISE-ROBUSTNESS ANALYSIS

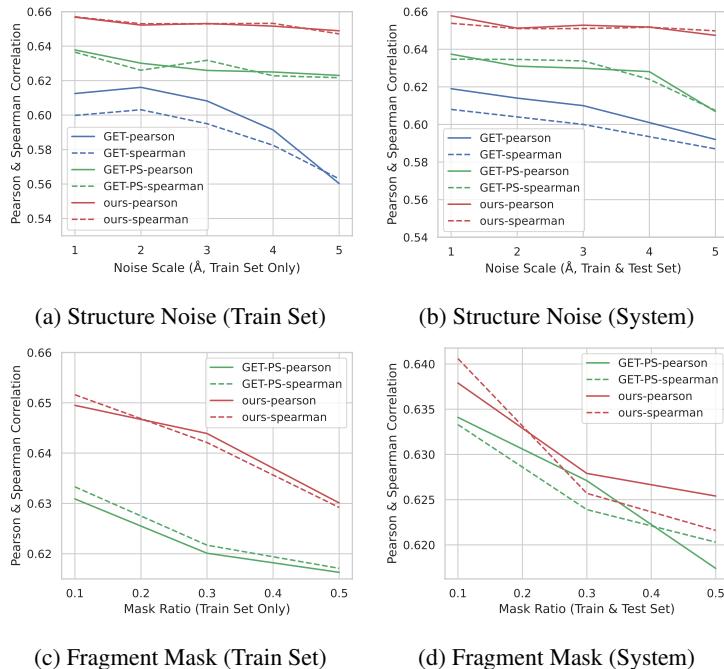
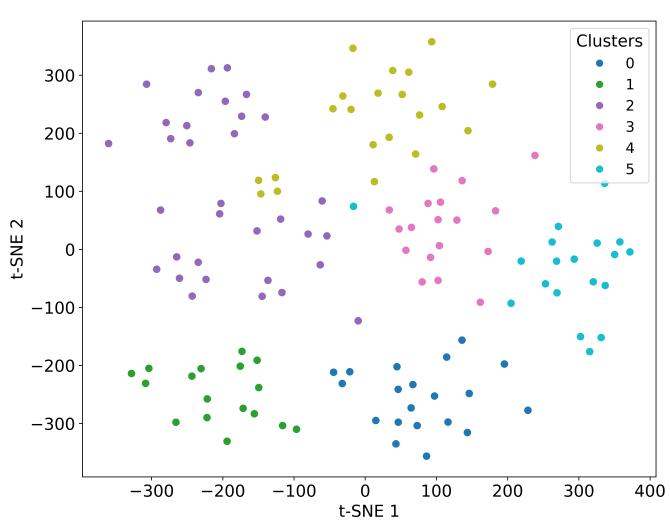


Figure 6: Noise robustness comparison on LBA. We report results from 3 runs.

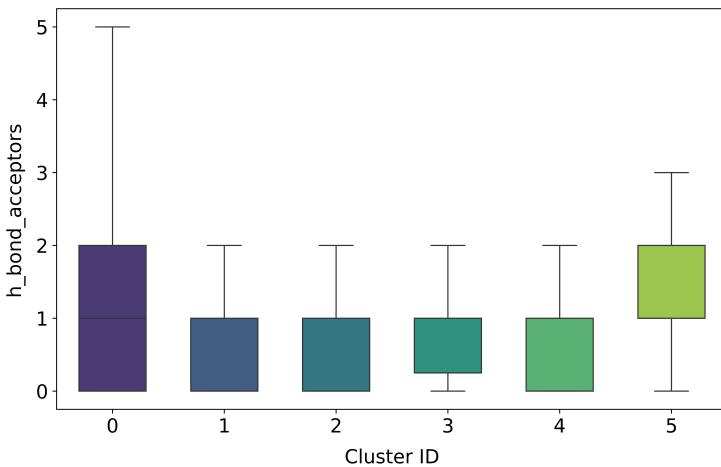
In this section, we analyze the robustness and generalization of GET, GET-PS, and our model under different noise scales on LBA. We consider two noise settings: adding noise only to the training set for simulating scenarios when training data is of low quality or is predicted, and adding noise to both training and test sets for simulating system bias or data resolution. To investigate the effect of noise in different features, we add random noise to input structures, or randomly mask token types with

Figure 7: **t-SNE Visualization of Fragment Embeddings.**

some ratio⁷. The results can be found in Figure 6. In general, we can conclude that: (i) Our model is robust to structure noise even up to 5 Å. (ii) Compared with atom-level 3D structures, token types provide stronger inductive bias for binding affinity prediction, which again highlights the importance of fragment-level representations.

D.5 FRAGMENT EMBEDDING ANALYSIS OF H-MINT

To evaluate the representation capability of the h-MINT model for molecular fragments, we first analyzed the spatial distribution of fragment embeddings using t-SNE, shown in Figure 7. The results showed that all fragments clustered into 6 distinct categories in the latent space.

Figure 8: **H-bond acceptors distribution across clusters.**

Statistical analysis of fragments in each cluster, as shown in Figure 8, revealed that Cluster 5 exhibited significant chemical specificity. This cluster was predominantly enriched with functional groups containing lone electron pairs on N and O atoms, such as C=O, CNC (C) =O, and CNS (=O) (=O).

⁷Since GET only uses atom representations for molecules, we only compare GET-PS with ours in (c)-(d).

1242 These structures, acting as typical hydrogen bond acceptors, can form stable interactions with polar
 1243 amino acids (e.g., the hydroxyl group of serine or the guanidinium group of arginine) in protein
 1244 binding pockets. This ability to capture key chemical features enables h-MINT to more sensitively
 1245 identify structural determinants affecting binding affinity, leading to superior performance over
 1246 existing sequence- or graph-based baseline models on the binding affinity prediction and virtual
 1247 screening tasks.

1248

1249 D.6 EFFECT OF THE FREQUENCY THRESHOLD

1250

1251 We vary the minimum frequency used by OverlapBPE, yielding vocabularies of sizes {98, 136, 212,
 1252 294}. Performance improves from small vocabularies to around 200, where we observe the best
 1253 aggregate results, and then slightly degrades at 294. Run-to-run variation is small ($\text{std} < 0.006$
 1254 across metrics), indicating stable behavior. These trends suggest a bias-variance trade-off: overly
 1255 strict thresholds (very small vocab) underfit by missing informative fragments, whereas overly lax
 1256 thresholds (very large vocab) admit rare or redundant fragments that increase sparsity and noise. A
 1257 moderate threshold around 200 offers the best balance between coverage and denoising.

1258

1259 Table 12: **Vocabulary size ablation on PDBBind.** Values are obtained over three runs.

1260

Vocab size	Pearson	Spearman	RMSE
98	0.618 ± 0.005	0.601 ± 0.006	1.344 ± 0.007
136	0.615 ± 0.002	0.605 ± 0.002	1.332 ± 0.003
212	0.640 ± 0.004	0.626 ± 0.005	1.299 ± 0.001
294	0.622 ± 0.002	0.617 ± 0.003	1.327 ± 0.002

1265

1266

1267 D.7 ADDITIONAL EXPERIMENTS ON MOLECULAR PROPERTY PREDICTION

1268

1269 We include 3 property prediction tasks from MoleculeNet. The baseline follows MoleculeNet directly,
 1270 which extracts ECFP features and trains XGBoost with grid search. The ECFP features are chosen
 1271 from 128-bit, 512-bit, 1024-bit and 2048-bit according to dataset. We augment the ECFP features
 1272 with bag-of-word features extracted from OverlapBPE and train the same XGBoost. The results are
 1273 in Table 13. The significant improvements in prediction error confirm that OverlapBPE provides
 1274 discriminative representations for molecules.

1275

1276 Table 13: Molecular Property Prediction Benchmarks from MoleculeNet.

1277

RMSE	ESOL ↓	FreeSolv ↓	Lipo ↓
ECFP	1.5668	3.9498	0.8875
ECFP + OverlapBPE	1.2972	3.3409	0.8270

1281

1282

1283 D.8 EXAMPLE TOKENS AND CHEMICAL INSIGHTS

1284

1285 By construction, OverlapBPE induces a hierarchical organization of fragments: basic tokens corre-
 1286 spond to chemically primitive units (atoms, bonds, individual rings), while composite tokens capture
 1287 larger patterns such as fused ring systems and side chains. In this section, we analyze how the learned
 1288 fragments align with standard functional chemistry.

1289

1290

1291

1292

1293

1294

1295

1296 We provide a list of the top 100 fragments mined from the LBA training set in Fig. 9, and we are
 1297 able to obtain many chemically and biologically meaningful subunits or motifs. We emphasize that
 1298 this motif-mining procedure is fully automatic, based solely on the data distribution, and leverages
 1299 no prior chemical or biological knowledge. We analyse and categorize the mined fragments into the
 1300 following chemically or biologically meaningful parts.

1301 **Chemical functional groups.** A functional group is a specific group of atoms or bonds within a
 1302 molecule that is responsible for its characteristic chemical properties and reactions. Representative
 1303 functional groups mined include: *carboxyl*, *phosphate*, *amide (peptide bond)*, *benzyl*, *secondary*

1296 *amino, tertiary amino, and quaternary amino (ammonium)*. These functional groups are small
 1297 chemical subunits that were often selected by hand in previous functional-group-based tokenization
 1298 but can be easily recovered with our approach.

1300 **Biomolecule subunits.** Compared to simple chemical functional groups, biomolecules such as
 1301 proteins, DNA, RNA, and polysaccharides are significantly larger. Yet, the monomers that constitute
 1302 these biomolecules exhibit characteristic patterns, such as amino acids, nucleotides, and saccharides.
 1303 Indeed, we observe a considerable number of biologically meaningful fragments in our codebook:
 1304 *peptide bond, adenine, pyranose, deoxyribose*, and, most notably, the whole nucleotide *adenosine*
 1305 monophosphate, which consists of an adenine, a ribose, and a phosphate. To the best of our
 1306 knowledge, none of the existing fragmentation approaches has been able to mine such large subunits
 1307 while preserving their biological significance.

1308 **Drug subunits.** Remarkably, we also observed that our approach could mine large chemical
 1309 subunits that are common motifs in small-molecule drugs. These include: *adamantane*, a 10-
 1310 carbon tricyclic motif with a highly symmetric fused four hexane rings, occurring in some antiviral
 1311 drugs; **sulfonamide**, in the antibacterial drug sulfanilamide; and most notably, the 28-heavy atom
 1312 **sulfonamide protease inhibitor motif** in HIV drugs. Indeed, our approach automatically mined this
 1313 large subunit, demonstrating its effectiveness at capturing the biochemical roles of many motifs.

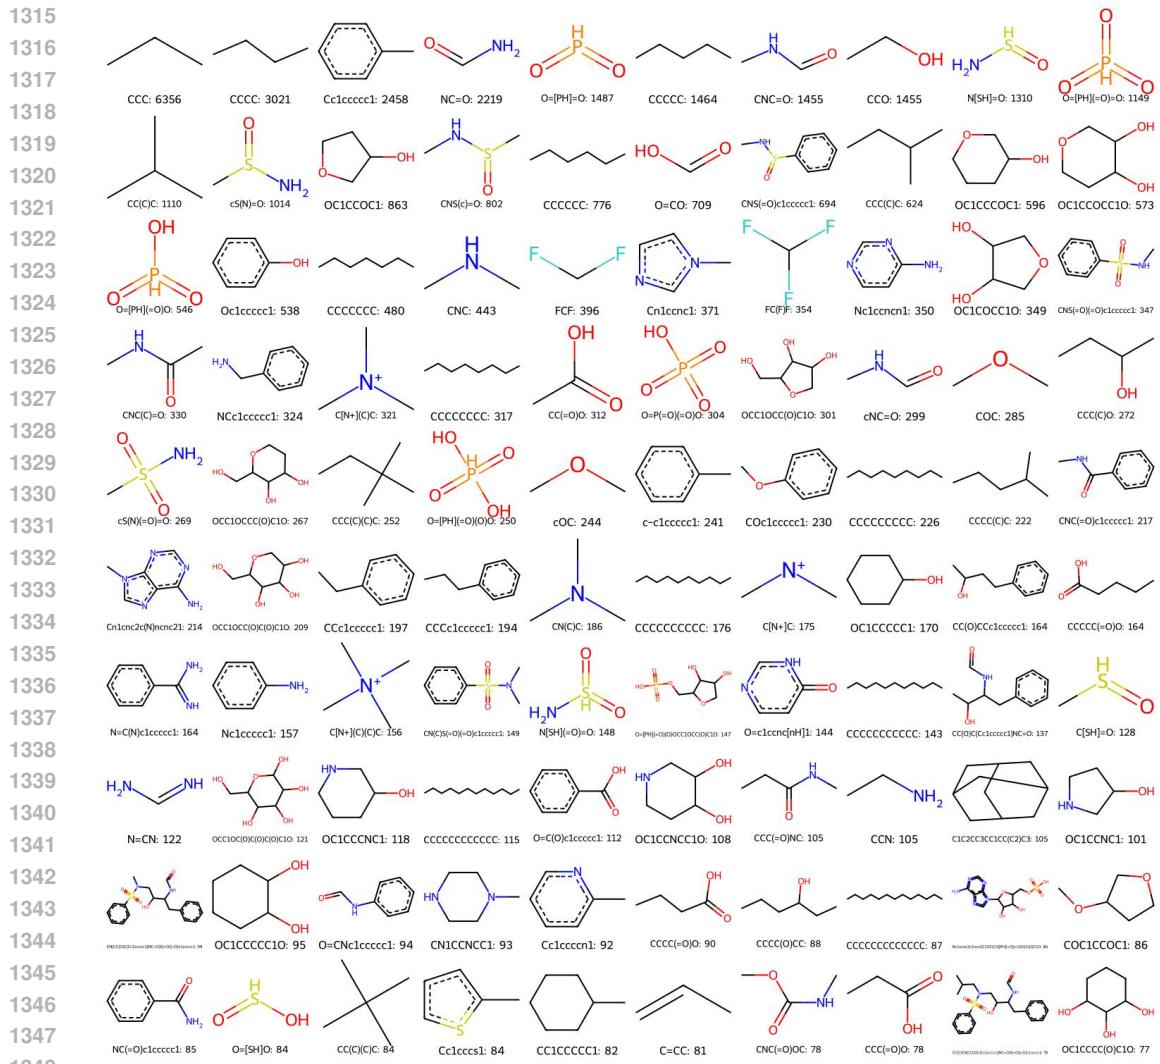


Figure 9: Top-100 composite tokens from LAB vocabulary.

1350 **E THE USE OF LARGE LANGUAGE MODELS**
13511352 We employ large language models exclusively for language editing, which is limited to polishing
1353 text to improve readability. No language models contributed to the development of research ideas,
1354 analysis, models, or interpretation of results.
1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403