

# GRAM-DTI: ADAPTIVE MULTIMODAL REPRESENTATION LEARNING FOR DRUG-TARGET INTERACTION PREDICTION

006 **Anonymous authors**

007 Paper under double-blind review

## ABSTRACT

013 Drug target interaction (DTI) prediction is a cornerstone of computational drug  
014 discovery, enabling rational design, repurposing, and mechanistic insights. While  
015 deep learning has advanced DTI modeling, existing approaches primarily rely  
016 on SMILES–protein pairs and fail to exploit the rich multimodal information  
017 available for small molecules and proteins. We introduce GRAM-DTI, a pre-  
018 training framework that integrates multimodal molecular and protein inputs into  
019 unified representations. GRAM-DTI extends volume-based contrastive learning  
020 to four modalities, capturing higher-order semantic alignment beyond conven-  
021 tional pairwise approaches. To handle modality informativeness, we propose  
022 adaptive modality dropout, dynamically regulating each modality’s contribution  
023 during pre-training. Additionally, IC50 activity measurements, when available,  
024 are incorporated as weak supervision to ground representations in biologically  
025 meaningful interaction strengths. Experiments on four publicly available datasets  
026 demonstrate that GRAM-DTI consistently outperforms state-of-the-art baselines.  
027 Our results highlight the benefits of higher-order multimodal alignment, adaptive  
028 modality utilization, and auxiliary supervision for robust and generalizable DTI  
029 prediction.

## 1 INTRODUCTION

030 Drug target interaction (DTI) prediction is a central challenge in computational drug discovery, un-  
031 derpinning applications in rational drug design, repurposing of approved drugs, and elucidation of  
032 mechanisms of action (Vefghi et al., 2025). Traditional experimental screening, though reliable, is  
033 prohibitively expensive and cannot feasibly cover the vast chemical and proteomic search space.  
034 Computational methods therefore play an increasingly critical role in prioritizing candidate drug-  
035 protein pairs for experimental validation, accelerating discovery pipelines and reducing cost (Panahandeh  
036 & Mansouri, 2025; Liao et al., 2025).

037 DTI prediction methods have evolved from similarity-based and network-based heuristics to ma-  
038 chine learning and, more recently, deep learning approaches (Shi et al., 2024; Panahandeh & Man-  
039 souri, 2025). Early methods relied on molecular similarity or interaction propagation but struggled  
040 with generalization. Modern neural models, including graph neural networks and sequence-based  
041 architectures now dominate, learning directly from raw SMILES and amino acid sequences (Peng  
042 et al., 2024; Zhao et al., 2025; Liu et al., 2025; Xia et al., 2023). However, these approaches remain  
043 largely restricted to SMILES–protein pairs, overlooking the richer multimodal information available  
044 for molecules and proteins that could yield more robust and generalizable interaction predictions.

045 While multimodal pre-training has been recently explored by few works for DTI prediction (Lu  
046 et al., 2025; Ye et al., 2021; Chen et al., 2020), existing approaches suffer from three limitations.  
047 Firstly, they rely on pairwise contrastive learning anchored to a single modality. Such schemes can-  
048 not capture higher-order interdependencies as the number of modalities increases (Cicchetti et al.,  
049 2024). Secondly, they assume all modalities are equally informative, ignoring that data sources  
050 often differ in quality, completeness, and relevance across samples and training stages. Static fu-  
051 sion can therefore lead to suboptimal representations when dominant but less informative modalities  
052 overshadow complementary signals. Finally, valuable supervision signals such as IC50 activity mea-  
053 surements are often missing or unreliable.

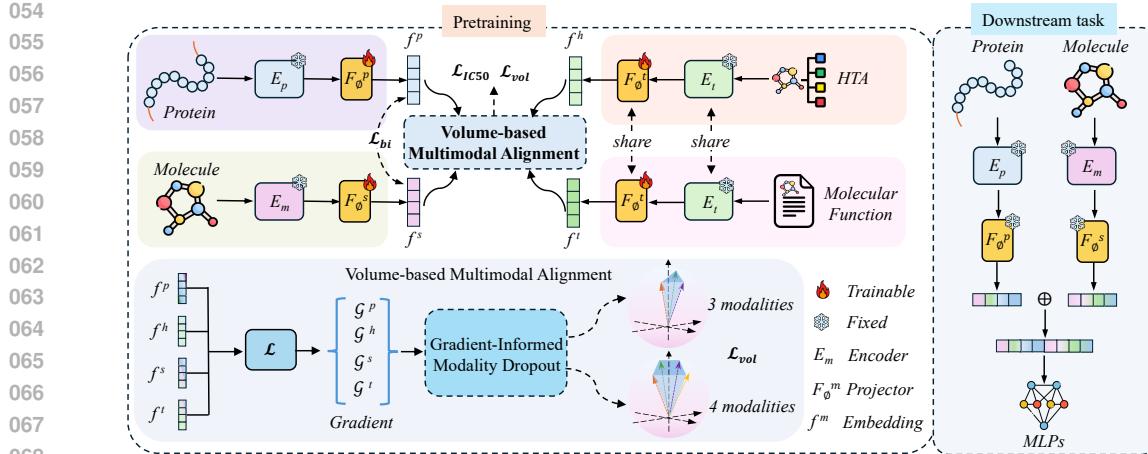


Figure 1: Overview of GRAM-DTI architecture. Left: pretraining phase with volume-based multimodal alignment across four modalities (SMILES, text, HTA, protein sequences). The framework uses gradient-informed adaptive modality selection to dynamically regulate modality contributions during training. Right: downstream task prediction.

surements are publicly available for a subset of drug–protein pairs, yet they remain unutilized during pre-training despite their direct biological relevance for DTI prediction task.

To address these gaps, we propose GRAM-DTI, a novel multimodal pre-training framework specifically tailored for downstream DTI prediction task (see Fig. 1). To this end, we curate a high quality multimodal dataset consisting of diverse protein and small molecule modalities and adapt the recent volume-based contrastive learning strategies from other domains (Cicchetti et al., 2024; Jiang et al., 2025a) for geometric alignment of these modalities. Unlike traditional contrastive learning techniques, this offers a theoretically principled and scalable approach for aligning multiple modalities. Beyond volume based contrastive learning, our framework is novel in its flexibility to learn to dynamically weight each modality based on its informativeness during pre-training while also supporting activity-based labels as auxiliary supervisory signals, when available. Our main contributions are as follows.

- We introduce GRAM-DTI, a pre-training framework for DTI that integrates multimodal small molecule protein modalities into a unified representation with volume-based contrastive learning.
- We introduce adaptive modality dropout, dynamically regulating modality contributions during pre-training to prevent dominant but less informative modalities from overwhelming complementary signals.
- We leverage IC50 activity measurements as additional weak auxiliary supervision, grounding learned representations in biologically meaningful drug–target interactions.
- We demonstrate state-of-the-art performance across four public datasets and multiple evaluation settings relevant for real-world drug discovery applications.

## 2 RELATED WORKS

**Multimodal Molecular Representation Learning** Recent advancements in molecular representation learning have shifted towards integrating multiple data modalities to enhance predictive performance. For instance, frameworks like TRIDENT (Jiang et al., 2025a) combine SMILES strings, hierarchical taxonomic annotations, and functional text of small molecules to capture richer molecular semantics. These approaches leverage contrastive learning pretraining to align diverse data sources, which improves generalization across various molecular downstream tasks even in the absence of fully paired datasets. Beyond TRIDENT, several molecular foundation models have been introduced, including MolFM (Luo et al., 2023) and MolCA (Liu et al., 2023), which integrate

108 molecular graphs, textual descriptions, and domain-specific annotations into unified representations.  
 109 These works highlight the broader trend of leveraging multimodal pre-training to construct general-  
 110 purpose molecular representations.  
 111

112 **Drug–Target Interaction (DTI) Prediction** DTI prediction has traditionally relied on unimodal  
 113 representations, such as SMILES strings for drugs and amino acid sequences for proteins. Early  
 114 deep learning models such as DeepDTA (Öztürk et al., 2018), MT-DTI (Shin et al., 2019), and  
 115 TransformerCPI (Chen et al., 2020) demonstrated the effectiveness of sequence-based architectures  
 116 for interaction prediction. Beyond sequence-based methods, more recent work has explored graph  
 117 neural networks and  $SE(3)$ -equivariant geometric deep learning models, such as GraphDTA (Nguyen  
 118 et al., 2021) and EquiBind (Stärk et al., 2022), which leverage spatial and structural information of  
 119 drugs and proteins to enhance binding affinity prediction. In parallel, knowledge graph–based meth-  
 120 ods such as NeoDTI (Wan et al., 2019) and Hetionet-based repurposing frameworks (Himmelstein  
 121 et al., 2017) exploit biomedical networks to capture higher-order relations among drugs, targets, and  
 122 diseases. More recently, multimodal approaches have been proposed to better capture the complex-  
 123 ity of drug–target interactions. For example, MDTips (Xia et al., 2023) integrates knowledge graphs,  
 124 gene expression profiles, and structural information, while MGNDTI (Peng et al., 2024) employs a  
 125 multimodal graph neural network to improve robustness and generalization. Another emerging di-  
 126 rection is pre-training with large-scale unlabeled data to mitigate the scarcity of labeled DTI pairs.  
 127 For instance, DTIAM (Lu et al., 2025) introduces separate pretraining for drug and target modalities  
 128 before merging the learned representations for DTI prediction.  
 129

130 **Modality Dropout** Modality dropout techniques have been proposed to enhance the robustness of  
 131 multimodal models by preventing over-reliance on any single modality. For instance, the Learnable  
 132 Irrelevant Modality Dropout (IMD) method (Alfasy et al., 2022) selectively drops irrelevant modal-  
 133 ities during training, improving performance in multimodal action recognition tasks. Additionally,  
 134 approaches like aggressive modality dropout have been shown to mitigate negative co-learning ef-  
 135 fects and enhance model accuracy in multimodal settings (Magal et al., 2025). Beyond dropout,  
 136 adaptive fusion mechanisms have also been investigated. Cross-attention and gating strategies (Tsai  
 137 et al., 2019; Peng et al., 2024; Mollaysa et al., 2025) dynamically regulate modality contributions,  
 138 while tensor fusion methods (Zadeh et al., 2017) capture higher-order interactions across modal-  
 139 ities. These ideas inform the design of adaptive strategies in molecular contexts, where modality  
 140 informativeness often varies across data sources and training stages.  
 141

142 Unlike existing works, our GRAM-DTI framework captures higher-order semantic relationships  
 143 beyond simple pairwise alignment/fusion. Furthermore, to the best of our knowledge, we are the  
 144 first to explore strategies for adaptive modality dropout in the context of DTI prediction.  
 145

### 3 METHODOLOGY

146 Building upon recent advances (Cicchetti et al., 2024; Jiang et al., 2025b) in volume-based modal-  
 147 ity alignment for effective representation learning, we extend the foundational concept of volume  
 148 loss (Cicchetti et al., 2024), originally formulated for audio-video-text data, to the domain of protein-  
 149 small molecule interactions. We aim to learn a unified embedding space that: 1) captures semantic  
 150 relationships across modalities; 2) remains robust when modalities vary in informativeness; and 3)  
 151 improves downstream DTI prediction task.  
 152

153 Formally, assume a pretraining dataset  $D = \{(x_i^s, x_i^t, x_i^h, x_i^p, \delta_{y_i}^{IC50})\}_{i=1}^N$ , where  $x_i^s$ ,  $x_i^t$ ,  $x_i^h$ , and  
 154  $x_i^p$  denote the SMILES sequence, textual description of molecule, hierarchical taxonomic anno-  
 155 tation (HTA) (Jiang et al., 2025b) of molecule, and protein sequence, respectively. The variable  
 156  $\delta_{y_i}^{IC50}$  indicates the IC50 activity class  $y_i^{IC50}$  if a measured IC50 value is available for the pro-  
 157 tein–molecule pair  $(x_i^p, x_i^s)$ , and 0 otherwise. As illustrated in Fig. 1, we employ pre-trained en-  
 158 coders  $E_i$  (MolFormer (Ross et al., 2022) for SMILES, MolT5 (Edwards et al., 2022) for text  
 159 and HTA, and ESM-2 (Lin et al., 2023) for proteins) to obtain initial modality-specific embed-  
 160 dings. To keep pre-training efficient and scalable, we freeze the backbone encoders and train  
 161 lightweight neural projectors  $F_\phi^m$  that map each modality embedding into a shared representation  
 162 space where they are semantically aligned. The resulting projected embeddings are denoted  $f^m$ ,  
 163 where  $m \in \{SMILES, text, HTA, protein\}$ .  
 164

162  
163

## 3.1 GRAMIAN VOLUME-BASED MULTIMODAL ALIGNMENT

164  
165  
166  
167

In contrast to traditional multimodal representation learning approaches which have been known to fail in capturing the complex interdependencies among three or more modalities (Cicchetti et al., 2024; Jiang et al., 2025b), volume loss uses Gramian volume-based alignment of modalities ensuring semantic coherence across all modalities simultaneously.

168  
169  
170  
171

**Gramian Volume** Given embeddings  $f_i^s, f_i^t, f_i^h, f_i^p \in \mathbb{R}^d$  that are learned from the four modalities  $x_i^s, x_i^t, x_i^h, x_i^p$  respectively, we first normalize them such that  $\|f_i^m\|_2 = 1$ . We can then construct the Gram matrix  $G \in \mathbb{R}^{4 \times 4}$  where

172  
173

$$G_{kj} = \langle f_i^k, f_i^j \rangle, \quad k, j \in \{s, t, h, p\} \quad (1)$$

174  
175  
176  
177

The 4-dimensional volume spanned by these embedded vectors is equal to the square root of the determinant of the Gramian matrix (Cicchetti et al., 2024):  $V(f_i^s, f_i^t, f_i^h, f_i^p) = \sqrt{\det(G)}$ . From multimodal alignment perspective, smaller volume intuitively suggests stronger semantic alignment, as the embeddings occupy a more compact and cohesive subspace and vice-versa.

178  
179  
180  
181

**Volume-Based Contrastive Loss** Given the Gramian volume, contrastive objective is cast as volume minimization/maximization. As proposed in (Cicchetti et al., 2024), to construct negative pairs, we chose an anchor modality  $a \in \{s, t, h, p\}$  as one of the four modalities. Therefore, for a batch of  $B$  samples, the contrastive loss on their learned embeddings is defined as follows:

182  
183  
184  
185

$$\mathcal{L}_{\text{vol}}^{\rightarrow} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(-V(a_i, f_i^t, f_i^h, f_i^p)/\tau)}{\sum_{j=1}^{B'} \exp(-V(a_j, f_i^t, f_i^h, f_i^p)/\tau)}, \quad (2)$$

186  
187  
188  
189  
190

where, for example, the first modality  $f_i^s$  is chosen as the anchor  $a_i$ , negative pairs are constructed by permuting the anchor, and  $\tau$  is the temperature parameter. We also add the reverse loss (w.r.t. negative pairs construction) to ensure symmetric alignment:  $\mathcal{L}_{\text{vol}}^{\leftarrow} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(-V(a_i, f_i^t, f_i^h, f_i^p)/\tau)}{\sum_{j=1}^{B'} \exp(-V(a_i, f_j^t, f_j^h, f_j^p)/\tau)}$ . The combined volume-based loss is

191  
192

$$\mathcal{L}_{\text{vol}} = \frac{1}{2} (\mathcal{L}_{\text{vol}}^{\rightarrow} + \mathcal{L}_{\text{vol}}^{\leftarrow}) \quad (3)$$

193

194

## 3.2 GRADIENT-INFORMED ADAPTIVE MODALITY SELECTION

195  
196  
197  
198  
199  
200  
201

While volume-based contrastive loss treats all modalities equally, different modalities may vary in quality and relevance, with contributions that change during training. Static fusion strategies risk either underutilizing weaker modalities or overfitting to dominant ones. We propose a gradient-informed modality dropout mechanism that dynamically adapts modality usage based on their instantaneous contribution to the loss function.

202  
203  
204

**Gradient Contribution Analysis** Assume  $\mathcal{L}_{\tilde{t}}$  denotes mini-batch loss at training step  $\tilde{t}$ . We measure the importance of modality  $m \in \{s, t, h, p\}$  by the magnitude of the gradient with respect to its embedding:

205  
206  
207

$$g_{\tilde{t}}^m = \left\| \frac{\partial \mathcal{L}_{\tilde{t}}}{\partial f_{\tilde{t}}^m} \right\|_2 \quad (4)$$

208  
209  
210  
211  
212

where  $f_{\tilde{t}}^m \in \mathbb{R}^d$  is the learned embedding of modality  $m$  at gradient step  $\tilde{t}$ . To avoid noisy decisions, we track the history of gradient contributions over the past  $K$  steps:  $\bar{g}_{\tilde{t}}^m = \frac{\sum_{k=0}^{K-1} \alpha^k g_{\tilde{t}-k}^m}{\sum_{k=0}^{K-1} \alpha^k}$ , where  $\alpha \in (0, 1)$  is an exponential decay factor which yields a smooth, temporally discounted importance score for each modality.

213  
214  
215

**Adaptive Modality Dropping Strategy** We employ a principled adaptive strategy that considers both the magnitude and variance of gradient contributions. Let  $\mu_{\tilde{t}} = \frac{1}{4} \sum_m \bar{g}_{\tilde{t}}^m$  and  $\sigma_{\tilde{t}} = \sqrt{\frac{1}{4} \sum_m (\bar{g}_{\tilde{t}}^m - \mu_{\tilde{t}})^2}$  denote the mean and standard deviation of weighted gradients across

216 modalities at the current gradient step  $\tilde{t}$ . We will drop a modality from the volume based contrastive  
 217 loss calculation with a probability of  $p_{\text{drop}}$ , which is a hyperparameter. The criteria to drop a modality  
 218 is defined as follows:

$$220 \quad m_{\text{drop}}^{(\tilde{t})} = \begin{cases} \arg \max_m \bar{g}_{\tilde{t}}^m & \text{if dominance detected, e.g., } \bar{g}_{\tilde{t}}^m > \mu_{\tilde{t}} + \lambda_{\sigma} \sigma_{\tilde{t}}, \\ 221 \quad \arg \min_m \bar{g}_{\tilde{t}}^m & \text{otherwise,} \\ 222 \quad \text{none} & \text{with probability } (1 - p_{\text{drop}}). \end{cases} \quad (5)$$

224 where  $\lambda_{\sigma} = 1.5$  is the threshold multiplier. This means that we adaptively drop modalities based on  
 225 two criteria: 1) *Dominance prevention*: if a modality’s contribution is much larger than others, we  
 226 drop it to avoid overfitting; 2) *Low-contribution pruning*: Otherwise, we drop the modality with the  
 227 smallest gradient contribution to encourage use of more informative signals. This dynamic selection  
 228 balances stability and diversity, ensuring all modalities remain engaged throughout training.

### 229 3.3 WEAK SUPERVISION THROUGH IC50 ACTIVITY MEASURE

231 As the IC50 values for wide range of protein-small molecule pairs are available on public data sources  
 232 such as BindingDB (Gilson et al., 2016), we introduce an additional classification task as an auxiliary  
 233 objective during pre-training. However, IC50 labels are not available for all possible protein-small  
 234 molecule pairs, this task provides only weak supervisory signal during pre-training when IC50 in-  
 235 formation is available. We train a classifier  $F_{\phi}^{IC50}$  to predict the IC50 class from the learned embed-  
 236 dings of all four modalities:  $f^{\text{fused}} = [f^s; f^t; f^h; f^p] \in \mathbb{R}^{4d}$ . Note that IC50 values are continuous,  
 237 but given the inherent challenges of IC50 regression, including heterogeneous value distributions,  
 238 wide dynamic ranges spanning several orders of magnitude, and noisy measurements (Qureshi et al.,  
 239 2015; Bavi et al., 2016; Ashraf et al., 2023), we formulate the problem as a three-class classification  
 240 task by employing discretizations on IC50 values (see Appendix A).

241 However, this discretizations comes with class-imbalance described in Appendix A. To address this  
 242 issue, we employ a weighted cross-entropy loss:

$$244 \quad \mathcal{L}_{\text{IC50}} = -\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} w_{y_i} \log p(y_i | f_i^{\text{fused}}), \quad (6)$$

246 where  $\mathcal{S}$  denotes the set of samples with valid IC50 annotations, and class weights are computed as:  
 247  $w_c = \frac{N_{\text{total}}}{C \cdot N_c}$ , where  $N_{\text{total}}$  being the total number of samples,  $C$  the number of classes, and  $N_c$  the  
 248 number of samples in class  $c$ .

250 **Auxiliary Bimodal Contrastive Loss** As the downstream task involves protein and molecule em-  
 251 beddings only, to emphasize alignment between these two, we also explicitly incorporate traditional  
 252 pairwise contrastive losses between SMILES and protein modalities:  $\mathcal{L}_{\text{bi}} = \frac{1}{2}(\mathcal{L}_{s \rightarrow p} + \mathcal{L}_{p \rightarrow s})$  where  
 253  $\mathcal{L}_{s \rightarrow p}$  and  $\mathcal{L}_{p \rightarrow s}$  follow the standard CLIP-style contrastive formulation (Radford et al., 2021).

### 255 3.4 UNIFIED TRAINING OBJECTIVE

257 The complete training objective integrates all components with appropriate weighting:

$$258 \quad \mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{vol}} + \lambda_2 \mathcal{L}_{\text{bi}} + \lambda_3 \mathcal{L}_{\text{IC50}} \quad (7)$$

260 where  $\lambda_1, \lambda_2, \lambda_3$  are hyperparameters. Note that  $\mathcal{L}_{\text{vol}}$  and  $\mathcal{L}_{\text{bi}}$  are applied on all the training instances  
 261 while  $\mathcal{L}_{\text{IC50}}$  are only applied for pairs of protein and molecule with valid IC50 annotations. For  
 262 gradient-based dropping of a modality in volume contrastive loss, we use  $\mathcal{L} = \lambda_2 \mathcal{L}_{\text{bi}} + \lambda_3 \mathcal{L}_{\text{IC50}}$ . See  
 263 Appendix C for details on model architecture and parameters.

## 264 4 EXPERIMENTS

### 266 4.1 DATASET

268 For pre-training, we employ the multimodal molecular dataset from TRIDENT (Jiang et al., 2025b),  
 269 consisting of 47,269 triplets of SMILES, text descriptions, and HTA annotations. We extend this

270 dataset by integrating protein binding information from BindingDB (Gilson et al., 2016), creating  
 271 quadruplets of  $\langle$ SMILES, Text, HTA, Protein $\rangle$  with IC50 measurements when available. To prevent  
 272 data leakage, we removed overlapping (SMILES, protein) pairs from our downstream evaluation  
 273 datasets. The final pretraining dataset contains 6,545 unique molecules and 4,418 proteins, forming  
 274 50,968 quadruplets, of which 16,035 include quantitative IC50 measurements for auxiliary super-  
 275 vision. See Appendix B for detailed dataset construction and statistics. **Ideally, we would remove**  
 276 **any drug or target that appears in the downstream tasks from the pretraining corpus.** However, given  
 277 the number of downstream tasks we evaluate, this would leave too little data for effective pretrain-  
 278 ing. Consequently, we only exclude overlaps at the (SMILES, protein) pair level. To verify that our  
 279 method does not memorize entity-specific patterns, we perform an overlap analysis on the Activation  
 280 dataset; results are provided in appendix E.3.

281 We evaluated our approach on four benchmark datasets from the DTIAM framework (Lu et al.,  
 282 2025). These datasets cover two types of prediction tasks: drug-target interaction (DTI) prediction  
 283 using the Yamanishi\_08 and Hetionet datasets, and mechanism of action (MoA) prediction using  
 284 the Activation and Inhibition datasets. **1) Activation dataset** obtained from the Therapeutic Target  
 285 Database (TTD) (Zhou et al., 2022), containing 1,426 drugs, 281 targets, and 1,913 known activation  
 286 interactions. **2) Yamanishi\_08** originally introduced by (Yamanishi et al., 2008) consists of four  
 287 sub-datasets: G-Protein Coupled Receptors, Ion Channels, Nuclear Receptors, and Enzymes. We  
 288 use the combined dataset constructed by (Ye et al., 2021), containing 791 drugs, 989 targets, and  
 289 5,127 known DTIs. **3) Hetionet dataset** constructed by (Himmelstein et al., 2017), which integrated  
 290 biomedical data from 29 public resources, comprising 1,384 drugs, 5,763 targets, and 49,942 DTIs.  
 291 **4) Inhibition dataset** derived from TTD (Zhou et al., 2022), containing 14,049 drugs, 1,088 targets,  
 292 and 21,055 known inhibition interactions. For detailed dataset statistics, see Appendix Table 3.

293 **Pre-training** Our four-modal contrastive learning framework employs a two-stage training  
 294 pipeline designed for computational efficiency and scalability. In the first stage, we extract em-  
 295 beddings using domain-specific pre-trained encoders: MoLFormer-XL (Ross et al., 2022) for  
 296 SMILES sequences, MolT5 (Edwards et al., 2022) for textual descriptions and HTA annotations,  
 297 and ESM2 (Lin et al., 2023) for protein sequences. In the second stage, we train lightweight pro-  
 298 jection networks that map these modality-specific embeddings into a unified representation space,  
 299 where volume-based contrastive alignment is performed using distributed training across multiple  
 300 GPUs. The complete training procedure, including our novel gradient-informed adaptive modality  
 301 dropout strategy, is detailed in Algorithms 1 and 2 in the Appendix.

302 Notably, we deliberately exclude  $\mathcal{L}_{\text{vol}}$  from the gradient computation for modality dropping. Instead,  
 303 we use  $\mathcal{L}_{\text{bi}}$  and  $\mathcal{L}_{\text{IC50}}$  to assess modality importance for two key reasons. First, the bimodal  
 304 contrastive loss and IC50 loss provide stable, interpretable signals about each modality’s contribution  
 305 without creating computational circularity. Second, IC50 values, though sparsely available, of-  
 306 fer biologically meaningful supervision that directly reflects protein-molecule interaction strength,  
 307 making the gradients from  $\mathcal{L}_{\text{IC50}}$  particularly valuable for identifying which modalities are most  
 308 important for drug-target activity prediction. Comprehensive training configuration details are pro-  
 309 vided in Appendix C. **In Table 4, we present the network architecture along with the hyperparameter**  
 310 **values used in our experiments. In Tables 6, 7, 8, and in Figure 4, we also provide the sensitivity**  
 311 **analysis with respect to the hyperparameters.**

312 We construct negative samples using an anchor point (Eq. 2); in each negative sample only a single  
 313 modality is altered while the remaining modalities remain aligned. We hypothesize this is the most  
 314 challenging negative sample scenario, since the model must distinguish the positive case, where all  
 315 modalities are aligned, from a negative case in which all but one modality are aligned. To assess  
 316 alternative strategies, we also evaluate an aggressive multi-domain negative sampling scheme in  
 317 which negative samples are formed by varying multiple modalities simultaneously. The results are  
 318 presented in Appendix section C.6.

319 **Downstream task** In the DTI and MoA prediction task, the objective is to determine whether a  
 320 given drug-target pair interacts, which constitutes a binary classification problem. Note that ex-  
 321 isting datasets only include those pairs that interact (positive class). Following standard practice  
 322 (Lu et al., 2025), we generated negative samples using a 1:10 ratio with positive samples for all  
 323 datasets. To evaluate the model’s generalization performance, we employed three different data  
 splitting strategies for train-test division: 1) *warm start*: The data is split based on protein-molecule

324  
325  
326  
327  
328Table 1: Mean performance comparison between GRAM-DTI and state-of-the-art baselines on DTI and MoA prediction tasks across multiple datasets and data splitting scenarios. GRAM-DTI demonstrates superior performance in most evaluation settings.  $\dagger$  indicates reproduced results; other results are from baseline papers. **Bold** denotes best performance.

Data	Metric	Scenario	CPL-GNN	MPNN-CNN	TransformerCPI	KGE-NFM	DTIAM $\dagger$	GRAM-DTI	Data	AI-DTI	DTIAM $\dagger$	GRAM-DTI
Yamanishi_08	AUPR	Warm start	0.431	0.816	0.802	0.817	0.901 $\pm$ 0.0085	<b>0.904<math>\pm</math>0.0079</b>	Activation	0.583	0.623 $\pm$ 0.0245	<b>0.642<math>\pm</math>0.0221</b>
		Drug cold start	0.167	0.408	0.410	0.341	0.439 $\pm$ 0.0580	<b>0.440<math>\pm</math>0.0662</b>		0.550	0.611 $\pm$ 0.0252	<b>0.628<math>\pm</math>0.0222</b>
		Target cold start	0.380	0.602	0.646	0.761	0.844 $\pm$ 0.0350	<b>0.849<math>\pm</math>0.0312</b>		0.219	0.391 $\pm$ 0.0320	<b>0.450<math>\pm</math>0.0374</b>
	AUROC	Warm start	0.821	0.952	0.953	0.948	0.967 $\pm$ 0.0050	<b>0.977<math>\pm</math>0.0042</b>		0.888	0.903 $\pm$ 0.0088	<b>0.914<math>\pm</math>0.0078</b>
		Drug cold start	0.629	0.797	0.767	0.779	0.818 $\pm$ 0.0255	<b>0.828<math>\pm</math>0.0285</b>		0.879	0.907 $\pm$ 0.0076	<b>0.913<math>\pm</math>0.0068</b>
		Target cold start	0.800	0.856	0.870	0.923	0.941 $\pm$ 0.0180	<b>0.955<math>\pm</math>0.0155</b>		0.652	0.792 $\pm$ 0.0240	<b>0.834<math>\pm</math>0.0258</b>
	Hetionet	Warm start	0.441	0.734	-	0.789	<b>0.879<math>\pm</math>0.0095</b>	0.859 $\pm$ 0.0082	Inhibition	0.840	<b>0.845<math>\pm</math>0.0070</b>	0.785 $\pm$ 0.0061
		Drug cold start	0.219	0.453	-	0.391	0.514 $\pm$ 0.0680	<b>0.529<math>\pm</math>0.0626</b>		<b>0.830</b>	0.731 $\pm$ 0.0045	0.756 $\pm$ 0.0034
		Target cold start	0.433	0.470	-	<b>0.651</b>	0.625 $\pm$ 0.0210	0.626 $\pm$ 0.0239		0.215	0.445 $\pm$ 0.0620	<b>0.464<math>\pm</math>0.0559</b>
	AUROC	Warm start	0.810	0.956	-	0.968	0.957 $\pm$ 0.0015	<b>0.981<math>\pm</math>0.0011</b>		0.952	<b>0.954<math>\pm</math>0.0025</b>	0.949 $\pm$ 0.0018
		Drug cold start	0.685	0.831	-	0.803	0.752 $\pm$ 0.0355	<b>0.855<math>\pm</math>0.0385</b>		<b>0.948</b>	0.921 $\pm$ 0.0028	0.940 $\pm$ 0.0018
		Target cold start	0.810	0.858	-	0.915	0.917 $\pm$ 0.0090	<b>0.921<math>\pm</math>0.0079</b>		0.605	0.819 $\pm$ 0.0205	<b>0.823<math>\pm</math>0.0028</b>

329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377

pairs, ensuring that no common pairs appear in both the training and test sets. 2) *drug cold start*: This split is performed at the molecule level, guaranteeing that no drug in the test set is present in the training set. 3) *target cold start*: Similar to the above, but split at the protein level, meaning no protein in the test set is seen during training. These three settings allow us to assess how well the model performs when faced with unseen molecule-protein pairs, unseen molecules, or unseen proteins, respectively. For evaluation, we followed the cross-validation protocols established in the original DTIAM framework (Lu et al., 2025): 10-fold cross-validation for DTI prediction tasks (Yamanishi\_08 and Hetionet datasets) and 5-fold cross-validation for MoA prediction tasks (Activation and Inhibition datasets). **Note that we generated negative samples at a 1:10 ratio relative to positive samples across all datasets, to ensure consistency with baseline methods and a fair comparison. Additional results with varying negative-sample ratios are provided in Appendix E.4, illustrating how our model performance changes as the ratio is adjusted.**

## 4.2 EXPERIMENTAL RESULTS

We evaluated GRAM-DTI against state-of-the-art models across multiple benchmark datasets to demonstrate its effectiveness. For DTI prediction tasks, Table 1 presents a comparison with five baselines: CPL-GNN (Tsubaki et al., 2019), MPNN-CNN (Gilmer et al., 2017), TransformerCPI (Chen et al., 2020), and KGE-NFM (Ye et al., 2021) and DTIAM (Lu et al., 2025), on the Yamanishi\_08 and Hetionet datasets. For MoA prediction tasks, we compared GRAM-DTI against two baselines: AI-DTI (Lee et al., 2023) and DTIAM (Lu et al., 2025) on the Activation and Inhibition datasets. The different baseline sets reflect the distinct methodological approaches and evaluation standards established for DTI and MoA prediction in the computational drug discovery community and follows prior works (Lu et al., 2025; Panahandeh & Mansouri, 2025).

GRAM-DTI demonstrates strong performance across benchmark datasets, with particularly notable gains in target cold start scenarios. For DTI tasks, our method achieves substantial improvements on Yamanishi\_08 in both warm start and target/drug cold start settings. On the larger Hetionet dataset, GRAM-DTI outperforms most baselines across multiple evaluation scenarios. For MoA prediction, GRAM-DTI consistently surpasses baselines on the Activation dataset, especially under target cold start conditions. On the Inhibition dataset, while GRAM-DTI does not outperform existing baselines in warm start and drug cold start settings, it exhibits excellent performance in target cold start.

Overall, GRAM-DTI outperforms state-of-the-art baselines in nearly all evaluation settings—10 out of 12 for DTI and 8 out of 12 for MoA tasks. Its strongest gains emerge on smaller datasets (Yamanishi\_08 and Activation), where pre-training provides the greatest benefit under limited supervision, thus validating its potential for real-world drug discovery applications with limited available labeled data. On larger datasets (Hetionet and Inhibition), GRAM-DTI remains on par with or outperforms strong baselines, particularly in cold start conditions. These results highlight the robustness and generalizability of our multimodal alignment framework, especially when extending to novel proteins.

## 4.3 ZERO-SHOT RETRIEVAL TASK

In addition to predicting drug–target interactions, an important aspect of evaluating our model’s effectiveness is its ability to accurately retrieve relevant molecules or proteins based on a given query. The retrieval task assesses the model’s capacity to learn meaningful, high-quality representations

378  
379  
380  
381  
382Table 2: Zero-shot retrieval performance comparison between GRAM-DTI and DTIAM baseline across four datasets. Results show Recall@K metrics for bidirectional retrieval tasks: S→P and P→S. GRAM-DTI demonstrates superior retrieval capability across most scenarios and datasets using only pretrained representations. **Bold** denotes best performance.

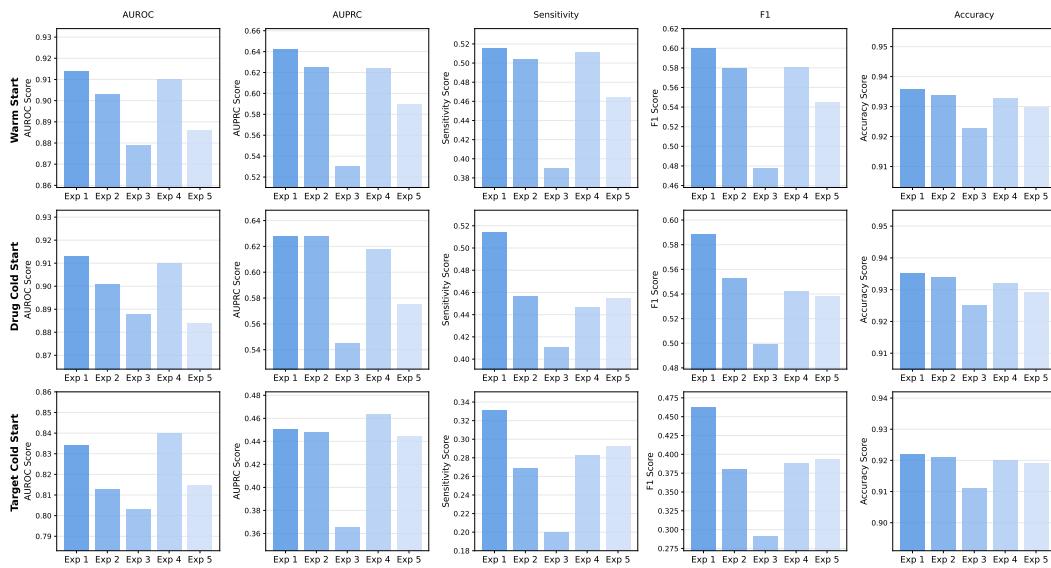
Direction	Metric	Yamanishi_08		Hetonet		Activation		Inhibition	
		DTIAM	GRAM-DTI	DTIAM	GRAM-DTI	DTIAM	GRAM-DTI	DTIAM	GRAM-DTI
S→P	R@1	0.0038±0.0004	<b>0.0465±0.0027</b>	0.0043±0.0002	<b>0.0331±0.0038</b>	0.0028±0.0002	<b>0.0136±0.0011</b>	0.0004±0.0000	<b>0.0055±0.0003</b>
	R@10	0.0341±0.0042	<b>0.1691±0.0084</b>	0.0434±0.0051	<b>0.1340±0.0025</b>	0.0266±0.0037	<b>0.1020±0.0067</b>	0.0097±0.0006	<b>0.0337±0.0011</b>
	R@100	0.1960±0.0181	<b>0.4449±0.0075</b>	0.2066±0.0109	<b>0.3616±0.0063</b>	0.3184±0.0229	<b>0.5688±0.0172</b>	0.1036±0.0104	<b>0.1994±0.0018</b>
P→S	R@1	0.0040±0.0002	<b>0.0742±0.0120</b>	<b>0.0404±0.0028</b>	0.0236±0.0010	0.0071±0.0008	<b>0.0370±0.0069</b>	0.0000±0.0000	<b>0.0221±0.0061</b>
	R@10	0.0849±0.0089	<b>0.2465±0.0256</b>	<b>0.1319±0.0095</b>	0.1049±0.0055	0.0463±0.0050	<b>0.2454±0.0142</b>	0.0028±0.0004	<b>0.0819±0.0065</b>
	R@100	0.3670±0.0186	<b>0.5540±0.0148</b>	0.3632±0.0474	<b>0.3841±0.0082</b>	0.2206±0.0264	<b>0.6029±0.0231</b>	0.0588±0.0049	<b>0.2325±0.0094</b>

383

389 that preserve semantic relationships across different modalities. This task is particularly relevant  
390 for applications such as drug repurposing and target identification (Luo et al., 2016; Pushpkom  
391 et al., 2019), where retrieving similar compounds or proteins can guide experimental validation and  
392 discovery.

393 To evaluate the retrieval capability of GRAM-DTI, we conduct a series of experiments across the  
394 same four datasets. For each dataset, we formulate two retrieval scenarios: (i) retrieving proteins  
395 given a drug query (S→P), and (ii) retrieving drugs given a protein query (P→S). Using the learned  
396 representations directly from our pre-training framework without any additional training, we com-  
397 pute similarity scores between query and candidate items. The performance is measured using  
398 standard metrics, including Recall@K (R@1, R@10, R@100), which indicate the proportion of  
399 relevant items retrieved within the top-K results.

400 The results, summarized in Table 2, demonstrate that our method outperforms DTIAM (the best  
401 baseline from DTI and MoA experiments in Table 1) across nearly all datasets and metrics. Not-  
402 ably, the superior performance in R@1 and R@10 indicates that our model effectively captures the  
403 semantic relationships necessary for accurate retrieval, highlighting the quality of the learned multi-  
404 modal representations. These strong zero-shot retrieval results provide compelling evidence that  
405 our multimodal pretraining framework successfully learns meaningful drug-target representations  
406 that generalize well beyond the specific downstream prediction tasks. More detailed experimental  
407 results can be found in Appendix E.

428  
429  
430  
431  
Figure 2: Ablation study results on the Activation dataset across five experimental configurations  
and three data splitting scenarios. The full GRAM-DTI model (Exp 1) outperforms variants with  
removed components in most cases, demonstrating the synergistic contribution of each training ob-  
jective component.

432 4.4 ABLATION STUDY  
433

434 Note that our main pre-training objective consists of three components (see Eq.7). To evaluate  
435 the contribution of each component, we conducted a comprehensive ablation study, comparing the  
436 performance of our model with each component systematically removed. We conduct five ablation  
437 experiments to evaluate the contribution of each component. **Exp 1** uses the full objective with  
438 modality dropout applied on volume loss calculation, i.e.,  $\mathcal{L} = \mathcal{L}_{\text{total}}$ , which is the same as our  
439 GRAM-DTI setup. **Exp 2** pre-trains without volume loss, using  $\mathcal{L} = \lambda_2 \mathcal{L}_{\text{bi}} + \lambda_3 \mathcal{L}_{\text{IC50}}$ . **Exp 3**  
440 pre-trains without traditional pairwise contrastive loss, employing  $\mathcal{L} = \lambda_1 \mathcal{L}_{\text{vol}} + \lambda_3 \mathcal{L}_{\text{IC50}}$ . **Exp 4**  
441 pre-trains without IC50 supervision, using  $\mathcal{L} = \lambda_1 \mathcal{L}_{\text{vol}} + \lambda_2 \mathcal{L}_{\text{bi}}$ . Finally, **Exp 5** uses the full objective  
442 but without modality dropout. The ablation study results on Activation dataset is presented in Figure  
443 2 while the same for Yamanishi\_08 dataset is reported in Appendix Figure 5. Across all setups, the  
444 full GRAM-DTI model (Exp 1) with all components enabled generally outperforms other variants  
445 where one component is removed.

446 **Impact of Gramian Volume-Based Alignment.** Gramian volume-based alignment provides sub-  
447 stantial benefits across most evaluation scenarios. Comparing it (Exp 1) with the variant excluding  
448 volume loss (Exp 2) reveals consistent improvements across the majority of metrics, particularly in  
449 challenging scenarios like target cold start where models must generalize to previously unseen pro-  
450 teins. The volume-based approach effectively captures higher-order relationships among the four  
451 modalities that cannot be achieved through pairwise alignments alone, leading to more robust mul-  
452 timodal representations.

453 **Impact of IC50 Auxiliary Supervision and Contrastive Loss.** Incorporating IC50 auxiliary su-  
454 pervision consistently improves performance across most evaluation scenarios (with the exception  
455 of Activation target cold start) as seen by comparing Exp 1 with Exp 4 (without IC50 supervision).  
456 Same conclusion holds when comparing Exp 1 with Exp 3, which suggests that the bimodal con-  
457 trastive loss also ensures robust drug-protein alignment and complements volume-based alignment.  
458 Together, these components capture both molecular activity principles and critical drug-protein re-  
459 lationships for effective prediction.

460 **Impact of Adaptive Modality Dropout.** Removing adaptive modality dropout (Exp 5), we see  
461 in figure 2, the performance consistently deteriorates, often by a large margin, compared to the  
462 no-dropout setting. By dynamically regulating modality contributions during training, the adaptive  
463 dropout prevents dominant modalities from overwhelming complementary signals while ensuring all  
464 modalities remain engaged. This prevents overfitting to specific modality combinations, ultimately  
465 leading to more generalizable representations. [To further validate this design choice, we compared](#)  
466 our probabilistic dropout strategy against "soft" weighting alternatives (e.g., weighted-modality  
467 gradients). Results (see Appendix E.6) demonstrate that our "hard" dropout strategy provides a stronger  
468 regularization effect and superior downstream performance.

469 **Impact of Molecular Encoder Strength.** To assess the modularity of our framework, we eval-  
470 uated GRAM-DTI using more advanced molecular encoders, specifically Uni-Mol2 (Ji et al., 2024)  
471 and BioT5+ (Pei et al., 2024), in place of MolFormer. As detailed in Appendix E.7, we observe that  
472 stronger encoders yield further performance gains, confirming that GRAM-DTI effectively leverages  
473 improvements in upstream foundation models.

474 **Multimodal Embedding Evolution** To visualize how GRAM-DTI learns unified representations,  
475 we examine embedding evolution across training epochs using t-SNE on 3,000 randomly sampled  
476 quadruplets (Figure 3). Initially, the four modalities form distinct, separate clusters. As training  
477 progresses, volume-based alignment gradually transforms rigid modality boundaries into semantically  
478 integrated representations while preserving modality-specific structures. By epoch 40, embeddings  
479 show substantial cross-modal integration where instances cluster by semantic relationships rather  
480 than purely by modality type. This evolution pattern provides visual evidence that our approach  
481 successfully balances cross-modal alignment with modality-specific information retention, support-  
482 ing the quantitative improvements observed in downstream tasks. Additional analyses with varying  
483 sample sizes are provided in Appendix F.1.

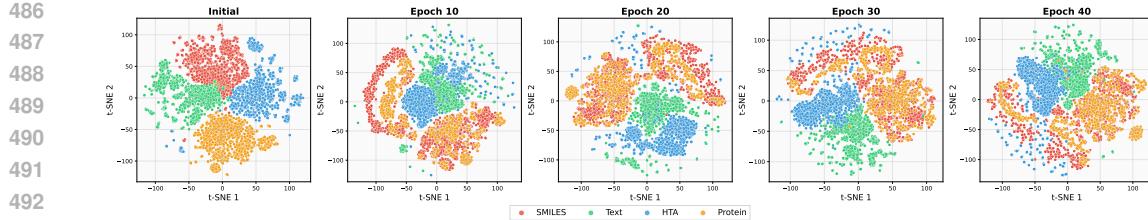


Figure 3: Evolution of multimodal embeddings during GRAM-DTI pre-training visualized using t-SNE on 3,000 samples. Four modalities (SMILES, Text, HTA, Protein) progressively align from separate clusters to semantically integrated representations, demonstrating effective volume-based multimodal alignment.

**Impact of missing modalities during pre-training** Bringing in as many relevant modalities in the pretraining would help learn better representation for the corresponding downstream task. However, what if during training, not all the modalities are available? We investigated this question by considering scenario where certain modality is not available during pretraining. The results are presented in Appendix section E.2. Moreover, we can extend pretraining to include samples with missing modalities, which would substantially increase the size of our training set. To assess whether all modalities are beneficial, our current pretraining phase includes only samples in which all four modalities are present, a choice that significantly limits the dataset. As a proof of concept, we evaluated whether including samples with only a subset of modalities improves downstream performance. The results, presented in Appendix E.5, indicate a promising direction: incorporating partial-modality samples can expand the pretraining corpus and may enhance model performance.

#### 4.5 FALSE NEGATIVE CASE ANALYSIS

To understand better when the model fails to predict drug target activity, we first systematically identified the top 10 "hardest" false negatives in the Activation dataset—pairs where the model predicted a strong negative signal despite a positive ground truth label. These are listed in table 22 in the appendix section E.8. From this list, we performed a detailed case study on Rank: Drug D03XIS (R-568) targeting T92076 (CASR). Our analysis suggests this prediction difficulty likely stems from the unique and complex biology of this pair, which is statistically rare in typical drug-target datasets: This case may be challenging because CASR is a Class C GPCR, fundamentally different from the Class A GPCRs that dominate drug databases. Three key factors may contribute to the prediction difficulty: (1) CASR has a large extracellular Venus flytrap domain, contrasting with the compact transmembrane binding pockets typical of Class A GPCRs; (2) it functions as an obligate homodimer with complex inter-protomer allosteric signaling; (3) R-568 acts as a positive allosteric modulator rather than a traditional orthosteric agonist. The prediction difficulty may reflect the biological rarity of Class C GPCR allosteric modulators in drug discovery.

## 5 CONCLUSION

We presented GRAM-DTI, a multimodal pretraining framework that extends volume-based contrastive learning to four modalities with gradient-informed adaptive modality dropout and IC50 auxiliary supervision. Evaluation across four benchmark datasets shows GRAM-DTI consistently outperforms baselines, particularly in cold start scenarios. Ablation studies (Appendix section 4.4) confirm synergistic contributions of each component. These results highlight the potential of multimodal pretraining for drug discovery, where integrating diverse data sources leads to more robust prediction models. Currently, the need to construct complete quadruplets  $\langle$ SMILES, Text, HTA, Protein $\rangle$  and remove overlapping (protein, SMILES) pairs with the downstream task has limited the scale of our pre-training dataset, restricting the diversity of molecules and proteins. To fully unlock the potential of GRAM-DTI and improve generalization to unseen molecular and protein targets, expanding the pre-training corpus will be crucial. In addition, incorporating protein-related modalities beyond sequence information could further enhance performance.

540 REFERENCES  
541

542 Saghir Alfasy, Jian Lu, Chen Xu, and Yuru Zou. Learnable irrelevant modality dropout for multi-  
543 modal action recognition on modality-specific annotated videos. In *Proceedings of the IEEE/CVF*  
544 *conference on computer vision and pattern recognition*, pp. 20208–20217, 2022.

545 Faisal Bin Ashraf, Sanjida Akter, Sumona Hoque Mumu, Muhammad Usama Islam, and Jasim  
546 Uddin. Bio-activity prediction of drug candidate compounds targeting sars-cov-2 using machine  
547 learning approaches. *Plos one*, 18(9):e0288053, 2023.

548 Rohit Bavi, Raj Kumar, Light Choi, and Keun Woo Lee. Exploration of novel inhibitors for bru-  
549 ton's tyrosine kinase by 3d qsar modeling and molecular dynamics simulation. *PloS one*, 11(1):  
550 e0147190, 2016.

551 Lifan Chen, Xiaoqin Tan, Dingyan Wang, Feisheng Zhong, Xiaohong Liu, Tianbiao Yang, Xiaomin  
552 Luo, Kaixian Chen, Hualiang Jiang, and Mingyue Zheng. Transformercipi: improving compound-  
553 protein interaction prediction by sequence-based deep learning with self-attention mechanism and  
554 label reversal experiments. *Bioinformatics*, 36(16):4406–4414, 2020.

555 Giordano Cicchetti, Eleonora Grassucci, Luigi Sigillo, and Danilo Comminiello. Gramian multi-  
556 modal representation learning and alignment. *arXiv preprint arXiv:2412.11959*, 2024.

557 Carl Edwards, Tuan Lai, Kevin Ros, Garrett Honke, Kyunghyun Cho, and Heng Ji. Translation  
558 between molecules and natural language. *arXiv preprint arXiv:2204.11817*, 2022.

559 Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural  
560 message passing for quantum chemistry. In *International conference on machine learning*, pp.  
561 1263–1272. Pmlr, 2017.

562 Michael K Gilson, Tiqing Liu, Michael Baitaluk, George Nicola, Linda Hwang, and Jenny Chong.  
563 Bindingdb in 2015: a public database for medicinal chemistry, computational chemistry and sys-  
564 tems pharmacology. *Nucleic acids research*, 44(D1):D1045–D1053, 2016.

565 Daniel S Himmelstein, Antoine Lizee, Christine Hessler, Leo Brueggeman, Sabrina L Chen, Dex-  
566 ter Hadley, Ari Green, Pouya Khankhanian, and Sergio E Baranzini. Systematic integration of  
567 biomedical knowledge prioritizes drugs for repurposing. *eLife*, 6:e26726, 2017.

568 Xiaohong Ji, Zhen Wang, Zhifeng Gao, Hang Zheng, Linfeng Zhang, Guolin Ke, et al. Uni-mol2:  
569 Exploring molecular pretraining model at scale. *arXiv preprint arXiv:2406.14969*, 2024.

570 Feng Jiang, Mangal Prakash, Hehuan Ma, Jianyuan Deng, Yuzhi Gao, Amina Mollaysa, Tommaso  
571 Mansi, Rui Liao, and Junzhou Huang. TRIDENT: Tri-modal molecular representation learning  
572 with taxonomic annotations and local correspondence. In *ICML 2025 Workshop on Multi-modal*  
573 *Foundation Models and Large Language Models for Life Sciences*, 2025a. URL <https://openreview.net/forum?id=SpoTt62oLY>.

574 Feng Jiang, Mangal Prakash, Hehuan Ma, Jianyuan Deng, Yuzhi Guo, Amina Mollaysa, Tommaso  
575 Mansi, Rui Liao, and Junzhou Huang. Trident: Tri-modal molecular representation learning with  
576 taxonomic annotations and local correspondence. *arXiv preprint arXiv:2506.21028*, 2025b.

577 Won-Yung Lee, Choong-Yeol Lee, and Chang-Eop Kim. Predicting activatory and inhibitory drug-  
578 target interactions based on structural compound representations and genetically perturbed trans-  
579 scriptomes. *PLoS One*, 18(4):e0282042, 2023.

580 Qian Liao, Yu Zhang, Ying Chu, Yi Ding, Zhen Liu, Xianyi Zhao, Yizheng Wang, Jie Wan, Yi-  
581 jie Ding, Prayag Tiwari, et al. Application of artificial intelligence in drug-target interactions  
582 prediction: a review. *npj biomedical innovations*, 2(1):1, 2025.

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

586 Sizhe Liu, Yuchen Liu, Haofeng Xu, Jun Xia, and Stan Z Li. Sp-dti: subpocket-informed transformer  
587 for drug–target interaction prediction. *Bioinformatics*, 41(3):btaf011, 2025.

594 Zhiyuan Liu, Sihang Li, Yanchen Luo, Hao Fei, Yixin Cao, Kenji Kawaguchi, Xiang Wang, and  
 595 Tat-Seng Chua. Molca: Molecular graph-language modeling with cross-modal projector and uni-  
 596 modal adapter. *arXiv preprint arXiv:2310.12798*, 2023.

597

598 Zhangli Lu, Guoqiang Song, Huimin Zhu, Chuqi Lei, Xinliang Sun, Kaili Wang, Libo Qin, Yafei  
 599 Chen, Jing Tang, and Min Li. Dtiam: a unified framework for predicting drug-target interactions,  
 600 binding affinities and drug mechanisms. *Nature Communications*, 16(1):2548, 2025.

601 Heng Luo, William Mattes, Donna L Mendrick, and Huixiao Hong. Molecular docking for identi-  
 602 fication of potential targets for drug repurposing. *Current topics in medicinal chemistry*, 16(30):  
 603 3636–3645, 2016.

604 Yizhen Luo, Kai Yang, Massimo Hong, Xing Yi Liu, and Zaiqing Nie. Molfm: A multimodal  
 605 molecular foundation model. *arXiv preprint arXiv:2307.09484*, 2023.

606

607 Nicholas Magal, Minh Tran, Riku Arakawa, and Suzanne Nie. Negative to positive co-learning with  
 608 aggressive modality dropout. *arXiv preprint arXiv:2501.00865*, 2025.

609 Amina Mollaysa, Artem Moskale, Pushpak Pati, Tommaso Mansi, Mangal Prakash, and Rui Liao.  
 610 Biolangfusion: Multimodal fusion of dna, mrna, and protein language models. *arXiv preprint  
 611 arXiv:2506.08936*, 2025.

612

613 Tin Nguyen, Hien Le, Timothy P Quinn, Thin Nguyen, Trung Le, and Svetha Venkatesh. Graphdta:  
 614 prediction of drug–target binding affinity using graph convolutional networks. *Bioinformatics*, 37  
 615 (8):1140–1147, 2021.

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

618

619 Fatemeh Panahandeh and Najme Mansouri. A comprehensive review of neural network-based ap-  
 620 proaches for drug–target interaction prediction. *Molecular Diversity*, pp. 1–48, 2025.

621 Qizhi Pei, Lijun Wu, Kaiyuan Gao, Xiaozhuan Liang, Yin Fang, Jinhua Zhu, Shufang Xie, Tao Qin,  
 622 and Rui Yan. Biot5+: Towards generalized biological understanding with iupac integration and  
 623 multi-task tuning. *arXiv preprint arXiv:2402.17810*, 2024.

624

625 Lihong Peng, Xin Liu, Min Chen, Wen Liao, Jiale Mao, and Liqian Zhou. Mgndti: a drug–target in-  
 626 teraction prediction framework based on multimodal representation learning and the gating mech-  
 627 anism. *Journal of Chemical Information and Modeling*, 64(16):6684–6698, 2024.

628

629 Sudeep Pushpakom, Francesco Iorio, Patrick A Evers, K Jane Escott, Shirley Hopper, Andrew Wells,  
 630 Andrew Doig, Tim Guilliams, Joanna Latimer, Christine McNamee, et al. Drug repurposing:  
 631 progress, challenges and recommendations. *Nature reviews Drug discovery*, 18(1):41–58, 2019.

632

633 Abid Qureshi, Himani Tandon, and Manoj Kumar. Avp-ic50pred: multiple machine learning  
 634 techniques-based prediction of peptide antiviral activity in terms of half maximal inhibitory con-  
 635 centration (ic50). *Peptide Science*, 104(6):753–763, 2015.

636

637 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
 638 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
 639 models from natural language supervision. In *International conference on machine learning*, pp.  
 640 8748–8763. PMLR, 2021.

641

642 Jerret Ross, Brian Belgodere, Vijil Chenthamarakshan, Inkit Padhi, Youssef Mroueh, and Payel Das.  
 643 Large-scale chemical language representations capture molecular structure and properties. *Nature  
 644 Machine Intelligence*, 4(12):1256–1264, 2022.

645

646 Wen Shi, Hong Yang, Linhai Xie, Xiao-Xia Yin, and Yanchun Zhang. A review of machine learning-  
 647 based methods for predicting drug–target interactions. *Health Information Science and Systems*,  
 648 12(1):30, 2024.

649

650 Bonggun Shin, Sanghyun Park, Kyungsook Kang, and Joyce C Ho. Self-attention based molecule  
 651 representation for predicting drug–target interaction. *Proceedings of Machine Learning Research*,  
 652 106:955–970, 2019.

648 Hannes Stärk, Octavian-Eugen Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola.  
 649 Equibind: Geometric deep learning for drug binding structure prediction. In *International Con-*  
 650 *ference on Machine Learning*, pp. 20503–20521. PMLR, 2022.

651

652 Yao-Hung Hubert Tsai, Shaojie Bai, Makoto Yamada, Louis-Philippe Morency, and Ruslan  
 653 Salakhutdinov. Multimodal transformer for unaligned multimodal language sequences. In *Pro-*  
 654 *ceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 6558–  
 655 6569, 2019.

656

657 Masashi Tsubaki, Kentaro Tomii, and Jun Sese. Compound–protein interaction prediction with end-  
 658 to-end learning of neural networks for graphs and sequences. *Bioinformatics*, 35(2):309–318,  
 659 2019.

660

661 Ali Vefghi, Zahed Rahmati, and Mohammad Akbari. Drug-target interaction/affinity prediction:  
 662 Deep learning models and advances review. *Computers in Biology and Medicine*, 196:110438,  
 663 2025.

664

665 Feng Wan, Lin Hong, An Xiao, Tong Jiang, and Jianyang Zeng. Neodti: neural integration of  
 666 neighbor information from a heterogeneous network for discovering new drug–target interactions.  
 667 In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1237–1245,  
 668 2019.

669

670 Xiaoqiong Xia, Chaoyu Zhu, Fan Zhong, and Lei Liu. Mdtips: a multimodal-data-based drug–  
 671 target interaction prediction system fusing knowledge, gene expression profile, and structural  
 672 data. *Bioinformatics*, 39(7):btad411, 2023.

673

674 Yoshihiro Yamanishi, Michihiro Araki, Alex Gutteridge, Wataru Honda, and Minoru Kanehisa. Pre-  
 675 diction of drug–target interaction networks from the integration of chemical and genomic spaces.  
 676 *Bioinformatics*, 24(13):i232–i240, 2008.

677

678 Qing Ye, Chang-Yu Hsieh, Ziyi Yang, Yu Kang, Jiming Chen, Dongsheng Cao, Shibo He, and  
 679 Tingjun Hou. A unified drug–target interaction prediction framework based on knowledge graph  
 680 and recommendation system. *Nature communications*, 12(1):6775, 2021.

681

682 Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Tensor  
 683 fusion network for multimodal sentiment analysis. In *Proceedings of the 2017 Conference on*  
*684 Empirical Methods in Natural Language Processing*, pp. 1103–1114, 2017.

685

686 Siqin Zhang, Kuo Yang, Zhenhong Liu, Xinxing Lai, Zhen Yang, Jianyang Zeng, and Shao Li.  
 687 Drugai: a multi-view deep learning model for predicting drug–target activating/inhibiting mech-  
 688 anisms. *Briefings in bioinformatics*, 24(1), 2023.

689

690 Yanpeng Zhao, Yuting Xing, Yixin Zhang, Yifei Wang, Mengxuan Wan, Duoyun Yi, Chengkun  
 691 Wu, Shangze Li, Huiyan Xu, Hongyang Zhang, et al. Evidential deep learning-based drug–target  
 692 interaction prediction. *Nature Communications*, 16(1):6915, 2025.

693

694

695 Ying Zhou, Yintao Zhang, Xichen Lian, Fengcheng Li, Chaoxin Wang, Feng Zhu, Yunqing Qiu, and  
 696 Yuzong Chen. Therapeutic target database update 2022: facilitating drug discovery with enriched  
 697 comparative data of targeted agents. *Nucleic acids research*, 50(D1):D1398–D1407, 2022.

698

699

700

701

## 702 A IC50 VALUES DISCRETIZATIONS

704 Given the inherent challenges of IC50 regression—including heterogeneous value distributions,  
 705 wide dynamic ranges spanning several orders of magnitude, and noisy measurements—we for-  
 706 mulate the problem as a three-class classification task. The IC50 values are discretized based on  
 707 pharmaceutical relevance thresholds:

$$709 \quad 710 \quad 711 \quad 712 \quad 713 \quad 714 \quad 715 \quad 716 \quad 717 \quad 718 \quad 719 \quad 720 \quad 721 \quad 722 \quad 723 \quad 724 \quad 725 \quad 726 \quad 727 \quad 728 \quad 729 \quad 730 \quad 731 \quad 732 \quad 733 \quad 734 \quad 735 \quad 736 \quad 737 \quad 738 \quad 739 \quad 740 \quad 741 \quad 742 \quad 743 \quad 744 \quad 745 \quad 746 \quad 747 \quad 748 \quad 749 \quad 750 \quad 751 \quad 752 \quad 753 \quad 754 \quad 755$$

$$\text{IC50 class} = \begin{cases} 0 & \text{if } \text{IC50} < 10\mu\text{M (effective)} \\ 1 & \text{if } 10\mu\text{M} \leq \text{IC50} \leq 1000\mu\text{M (moderate)} \\ 2 & \text{if } \text{IC50} > 1000\mu\text{M (ineffective)} \end{cases} \quad (8)$$

This discretization strategy aligns with established drug discovery practices (Qureshi et al., 2015; Bavi et al., 2016; Ashraf et al., 2023) where compounds with  $\text{IC50} < 10\mu\text{M}$  are considered highly active, those between  $10 - 1000\mu\text{M}$  show moderate activity, and those  $> 1000\mu\text{M}$  are typically considered inactive.

## B DATASET

**720 Pretraining Data** Our pretraining dataset builds upon the high-quality multimodal molecular  
 721 dataset from TRIDENT (Jiang et al., 2025b), which provides comprehensive molecular repres-  
 722 entations through the integration of SMILES strings, natural language descriptions, and Hierarchical  
 723 Taxonomic Annotations (HTA). The original TRIDENT dataset contains 47,269 carefully curated  
 724  $\langle \text{SMILES}, \text{Text}, \text{HTA} \rangle$  triplets sourced from PubChem, where each molecule is annotated across 32  
 725 diverse taxonomic classification systems.

To enable protein-molecule interaction modeling, we extended this dataset by incorporating binding  
 727 affinity information from BindingDB, a comprehensive database of measured binding affinities for  
 728 protein-molecule interactions. We mapped molecules from the TRIDENT dataset to BindingDB  
 729 entries using molecular identifiers, creating 5-tuples of the form  $\langle \text{SMILES}, \text{Text}, \text{HTA}, \text{Protein},$   
 730  $\text{IC50} \rangle$ . This integration combines the rich semantic and structural information from TRIDENT with  
 731 quantitative binding affinity measurements, providing a unified multimodal representation that cap-  
 732 tures both molecular properties and protein-molecule interactions. Following standard practices in  
 733 molecular property prediction, we implemented careful data filtering to prevent information leak-  
 734 age between pretraining and downstream evaluation. Specifically, we removed all SMILES-protein  
 735 binding pairs that appear in our downstream task datasets to ensure fair evaluation and prevent over-  
 736 fitting to specific molecular-protein combinations seen during pretraining.

After filtering, 6,545 unique molecules have associated protein binding information. Considering  
 737 that each molecule can interact with multiple proteins, this results in a total of 50,968 quadruplets  
 738  $\langle \text{Protein}, \text{SMILES}, \text{Text}, \text{HTA} \rangle$ , covering 4,418 unique proteins. Among these quadruplets, 16,035  
 739 entries include quantitative IC50 measurements, providing high-quality binding affinity annotations  
 740 for modeling.

**742 Downstream Task Datasets** We evaluated our approach on four benchmark datasets (see Table 3)  
 743 from the DTIAM framework (Lu et al., 2025), covering drug-target interaction (DTI) prediction and  
 744 mechanism of action (MoA) prediction tasks. 1) **745 Activation dataset** obtained from the Therapeutic  
 746 Target Database (TTD) (Zhou et al., 2022), containing 1,426 drugs, 281 targets, and 1,913 known  
 747 activation interactions. 2) **Yamanishi\_08** originally introduced by (Yamanishi et al., 2008) and  
 748 consists of four sub-datasets: G-Protein Coupled Receptors (GPCR), Ion Channels (IC), Nuclear  
 749 Receptors (NR), and Enzymes (E). We use the combined dataset constructed by (Ye et al., 2021),  
 750 containing 791 drugs, 989 targets, and 5,127 known DTIs. 3) **Hetionet dataset** constructed by  
 751 (Himmelstein et al., 2017), which integrated biomedical data from 29 public resources, comprising  
 752 1,384 drugs, 5,763 targets, and 49,942 DTIs. 4) **Inhibition dataset** also derived from TTD (Zhou  
 753 et al., 2022), containing 14,049 drugs, 1,088 targets, and 21,055 known inhibition interactions.

The MoA refers to how a drug works on its target to produce the desired effects, which involve  
 754 two major roles: activation and inhibition mechanisms. Distinguishing the activation and inhibition  
 755 MoA between drugs and targets is critical and challenging in the drug discovery and development  
 756 process, as well as their clinical applications Zhang et al. (2023).

756  
 757 Table 3: Statistics of downstream task datasets for binary classification. Known Interactions repre-  
 758 sents the number of positive drug-target binding pairs, while Total Samples includes both positive  
 759 samples and 10 times negative samples generated following standard practice.

760 <b>Dataset</b>	761 <b>Task Type</b>	762 <b>Drugs</b>	763 <b>Targets</b>	764 <b>Known Interactions</b>	765 <b>Total Samples</b>
Yamanishi_08	DTI	791	989	5,127	56,397
Hetionet	DTI	1,384	5,763	49,942	549,362
Activation	MoA	1,426	281	1,913	21,043
Inhibition	MoA	14,049	1,088	21,055	231,605

## 766      C PRE-TRAINING SETUP AND ARCHITECTURAL DETAILS

### 767      C.1 PRE-TRAINING INFRASTRUCTURE

768      Our four-modal contrastive learning framework employs a two-stage training pipeline. First, we  
 769 extract embeddings from domain-specific pre-trained models (MoLFormer-XL (Ross et al., 2022)  
 770 for SMILES, MolT5(Edwards et al., 2022) for text/HTA, ESM2 (Lin et al., 2023) for proteins).  
 771 Second, we train projection networks and the GRAM4Modal loss using distributed training across  
 772 multiple GPUs. The complete training procedure is detailed in Algorithm 1, which incorporates our  
 773 gradient-based modality dropping strategy (Algorithm 2).

774      Notably, we deliberately exclude  $\mathcal{L}_{\text{vol}}$  from the gradient computation for modality dropping to avoid  
 775 circular dependency, where the volume loss computation would depend on gradients derived from  
 776 that same computation. Instead, we use  $\mathcal{L} = \lambda_2 \mathcal{L}_{\text{bi}} + \lambda_3 \mathcal{L}_{\text{IC50}}$  to assess modality importance for two  
 777 key reasons: 1) *Avoiding circular dependency*: The bimodal contrastive loss and IC50 loss provide  
 778 stable, interpretable signals about each modality’s contribution without creating computational cir-  
 779 cularity; 2) *Leveraging weak supervision*: IC50 values, though sparsely available, offer biologically  
 780 meaningful supervision that directly reflects protein-molecule interaction strength. The gradients  
 781 from  $\mathcal{L}_{\text{IC50}}$  thus provide valuable information about which modalities are most important for pre-  
 782 dicting drug-target activity, making them suitable signals for adaptive modality selection. Table 4  
 783 provides comprehensive training configuration details.

---

#### 792      **Algorithm 1** Four-Modal Contrastive Learning with Gradient-based Modality Dropping

---

793      **Require:** Pre-computed embeddings  $\{x_i^s, x_i^t, x_i^h, x_i^p\}$

794      **Require:** Drop probability  $p_{\text{drop}}$ , temperature  $\tau$

795      **Ensure:** Projected features  $\{f^s, f^t, f^h, f^p\}$

796      1:  $f^m \leftarrow F_\phi^m(E_m(x^m))$  for  $m \in \{s, t, h, p\}$

797      2:  $f^m \leftarrow \|f^m\|_2 = 1$  for all modalities

798      3:  $d \leftarrow \text{GradientBasedDrop}(\{f^m\}, \mathcal{L}, p_{\text{drop}})$

799      4: **if**  $d.\text{should\_drop} = \text{False}$  **then**

800       5:  $V_f \leftarrow \text{GRAM4Modal}(f^p, \{f_{\text{all}}^s, f_{\text{all}}^t, f_{\text{all}}^h\})$

801       6:  $V_r \leftarrow \text{GRAM4Modal}(f_{\text{all}}^p, \{f^s, f^t, f^h\})^T$

802       7: **else**

803       8:  $m_a \leftarrow d.\text{anchor\_modality}$

804       9:  $\{m_1, m_2\} \leftarrow \text{remaining\_modalities} \setminus \{m_a\}$

805       10:  $V_f \leftarrow \text{GRAM3Modal}(f^{m_a}, \{f_{\text{all}}^{m_1}, f_{\text{all}}^{m_2}\})$

806       11:  $V_r \leftarrow \text{GRAM3Modal}(f_{\text{all}}^{m_a}, \{f^{m_1}, f^{m_2}\})^T$

807       12: **end if**

808       13:  $S_f \leftarrow -V_f / \tau, S_r \leftarrow -V_r / \tau$

809       14:  $\mathcal{L}_{\text{vol}} \leftarrow \frac{1}{2} [\mathcal{L}_{\text{vol}}^{\rightarrow} + \mathcal{L}_{\text{vol}}^{\leftarrow}]$

15: **return**  $\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{vol}} + \lambda_2 \mathcal{L}_{\text{bi}} + \lambda_3 \mathcal{L}_{\text{IC50}}$

---

810 C.2 MODEL ARCHITECTURE  
811

812 The projection networks  $F_\phi^m$  map pre-computed embeddings to a unified 512-dimensional space.  
 813 Each projection consists of three linear layers with GELU activations, layer normalization, and  
 814 dropout (rate=0.1). The IC50 classification head  $F_\phi^{IC50}$  concatenates all four modality features  
 815  $f_{\text{fused}} = [f^s; f^t; f^h; f^p]$  and predicts binding affinity classes through a two-layer MLP with dropout  
 816 (rate=0.3). The pre-trained encoder specifications are detailed in Table 5. All encoders  $E_m$  are  
 817 frozen during training to leverage their pre-trained representations while only fine-tuning the pro-  
 818 jection networks  $F_\phi^m$  for computational efficiency.

820  
821 Table 4: Training Configuration Parameters

822 Parameter	823 Configuration
823 Hardware	824 Multi-GPU NVIDIA (CUDA)
824 Training framework	825 PyTorch DDP, NCCL
825 Batch size	826 1280 per GPU
826 Learning rate	827 $1 \times 10^{-4}$ (Adam)
827 Epochs	828 40
828 Temperature $\tau$	829 0.07
829 Drop probability $p_{\text{drop}}$	830 0.8
830 Gradient history length $K$	831 5
831 Decay factor $\alpha$	832 0.9
832 Threshold multiplier $\lambda_\sigma$	833 1.5
833 Loss weights $\lambda_1, \lambda_2, \lambda_3$	1.0, 1.0, 1.0
	Label smoothing 0.1

834  
835  
836  
837 **Algorithm 2** Gradient-based Adaptive Modality Dropping

838 **Require:** Features  $\{f^m\}_{m \in \{s, t, h, p\}}$ , current loss  $\mathcal{L}_{\tilde{t}}$ , drop probability  $p_{\text{drop}}$   
 839 **Require:** Gradient history length  $K$ , decay factor  $\alpha$ , threshold  $\lambda_\sigma = 1.5$

840 **Ensure:** Drop decision  $\{\text{should\_drop}, m_{\text{drop}}, \text{anchor\_modality}\}$

841 1: **if**  $\text{random}() > p_{\text{drop}}$  or not training **then**  
 842 2:   **return** {False, none, protein}  
 843 3: **end if**  
 844 4: **for**  $m \in \{s, t, h, p\}$  **do**  
 845 5:    $g_{\tilde{t}}^m \leftarrow \left\| \frac{\partial \mathcal{L}_{\tilde{t}}}{\partial f_{\tilde{t}}^m} \right\|_2$   
 846 6:   Update gradient history for modality  $m$   
 847 7: **end for**  
 848 8: **for**  $m \in \{s, t, h, p\}$  **do**  
 849 9:    $\bar{g}_{\tilde{t}}^m \leftarrow \frac{\sum_{k=0}^{K-1} \alpha^k g_{\tilde{t}-k}^m}{\sum_{k=0}^{K-1} \alpha^k}$   
 850 10: **end for**  
 851 11:  $\mu_{\tilde{t}} \leftarrow \frac{1}{4} \sum_m \bar{g}_{\tilde{t}}^m, \sigma_{\tilde{t}} \leftarrow \sqrt{\frac{1}{4} \sum_m (\bar{g}_{\tilde{t}}^m - \mu_{\tilde{t}})^2}$   
 852 12: **for**  $m \in \{s, t, h, p\}$  **do**  
 853 13:   **if**  $\bar{g}_{\tilde{t}}^m > \mu_{\tilde{t}} + \lambda_\sigma \sigma_{\tilde{t}}$  **then**  
 854 14:      $m_{\text{drop}}^{(\tilde{t})} \leftarrow m$ ; **break**  
 855 15:   **end if**  
 856 16: **end for**  
 857 17: **if**  $m_{\text{drop}}^{(\tilde{t})}$  not found **then**  
 858 18:    $m_{\text{drop}}^{(\tilde{t})} \leftarrow \arg \min_m \bar{g}_{\tilde{t}}^m$   
 859 19: **end if**  
 860 20:  $m_{\text{anchor}} \leftarrow \text{random\_choice}(\{s, t, h, p\} \setminus \{m_{\text{drop}}^{(\tilde{t})}\})$   
 861 21: **return** {True,  $m_{\text{drop}}^{(\tilde{t})}$ ,  $m_{\text{anchor}}$ }

864

865

Table 5: Pre-trained Encoder Specifications

866

867

Modality	Model $E_m$	Output Dim
SMILES ( $x^s$ )	MoLFormer-XL-both-10pct	768
Text ( $x^t$ )	MolT5-base	768
HTA ( $x^h$ )	MolT5-base (shared)	768
Protein ( $x^p$ )	ESM2_t33_650M_UR50D	1280

870

871

### C.3 COMPUTATIONAL EFFICIENCY

873

Our method is highly efficient because we freeze the large encoder backbones (ESM2, MolFormer, MolT5) and only train the lightweight projection layers. This significantly reduces computational and memory overhead.

877

- **Hardware:** All experiments were conducted on a single A6000 GPU.
- **Peak Memory:** The peak GPU memory usage during pretraining is only 0.12 GB\*.
- **Batch Size:** We use a large batch size of 1280.
- **Pretraining Speed:** Each pretraining epoch takes approximately \*\*3 seconds\*\* to complete.

883

This demonstrates that our method is not memory-intensive and is computationally very efficient.

884

885

886

### C.4 HYPERPARAMETER TUNING AND SENSITIVITY ANALYSIS

887

We tuned the hyperparameters for the pretrained model and eventually set the final optimal values as:  $\lambda_\sigma = 1.5$ ,  $\lambda_1 = 1$ ,  $\lambda_2 = 1$ ,  $\lambda_3 = 1$ ,  $K = 5$  as presented in table 4. To further analysis the model performance sensitivity with respect to each parameters, we investigated the impact of the each parameters value on the downstream task. In table 6, 7 and 8 we show the result on the Activation dataset when we change the parameters value from the optimal values that are used in our final model. These results demonstrate that our model’s performance is stable within a reasonable range of these hyperparameters, with our chosen settings providing a robust and effective performance.

895

896

Table 6: Downstream-task performance on the Activation dataset: sensitivity to changes in hyperparameter values from the optimal setting, evaluated under the Warm-start setup. The row in bold indicates the optimal hyperparameter values used during pretraining and the corresponding downstream performance.

900

901

$\lambda_1$	$\lambda_2$	$\lambda_3$	gradient_std_multiplier	$\lambda_\sigma$	gradient_history_length	$K$	AUPRC	AUROC
<b>1</b>	<b>1</b>	<b>1</b>		<b>1.5</b>		<b>5</b>	<b>0.6424</b>	<b>0.9142</b>
0.5	1	1		1.5		5	0.6449	0.9125
1	0.5	1		1.5		5	0.6237	0.9175
1	1	0.5		1.5		5	0.6326	0.9102
1	1	1		2		5	0.6340	0.9014
1	1	1		1.5		10	0.6155	0.9130

907

908

909

Table 7: Downstream-task performance on the Activation dataset: sensitivity to changes in hyperparameter values from the optimal setting, evaluated under the Drug cold start setup.

910

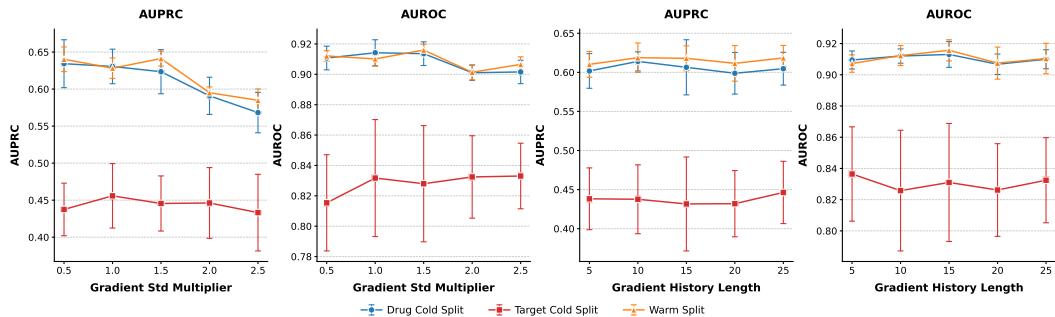
911

912

$\lambda_1$	$\lambda_2$	$\lambda_3$	gradient_std_multiplier	gradient_history_length	AUPRC	AUROC
<b>1</b>	<b>1</b>	<b>1</b>	<b>1.5</b>	<b>5</b>	<b>0.6278</b>	<b>0.9125</b>
0.5	1	1	1.5	5	0.6404	0.9129
1	0.5	1	1.5	5	0.6299	0.9107
1	1	0.5	1.5	5	0.6152	0.9114
1	1	1	2	5	0.6183	0.9009
1	1	1	1.5	10	0.6156	0.9103

918  
 919 **Table 8: Downstream-task performance on the Activation dataset: sensitivity to changes in hyper-**  
 920 **parameter values from the optimal setting, evaluated under the Target cold start setup.**

$\lambda_1$	$\lambda_2$	$\lambda_3$	gradient_std_multiplier	gradient.history_length	AUPRC	AUROC
<b>1</b>	<b>1</b>	<b>1</b>	<b>1.5</b>	<b>5</b>	<b>0.4497</b>	<b>0.8335</b>
0.5	1	1	1.5	5	0.4394	0.8224
1	0.5	1	1.5	5	0.4274	0.8190
1	1	0.5	1.5	5	0.4571	0.8273
1	1	1	2	5	0.4286	0.8270
1	1	1	1.5	10	0.4343	0.8295



928  
 929 **Figure 4: Effect of gradient optimization parameters on model performance. AUPRC and AUROC**  
 930 **scores are shown for varying gradient standard deviation multiplier (left) and gradient history length**  
 931 **(right) across three data split scenarios: drug cold split, protein (drug) cold split, and warm split.**  
 932 **Error bars indicate standard deviation.**

933  
 934 To further demonstrate the impact of the gradient\_std\_multiplier and gradient\_history\_length parameters  
 935 on model performance, we conducted a sensitivity analysis. We fixed all other parameters  
 936 and varied either gradient\_std\_multiplier or gradient\_history\_length to observe how performance  
 937 changed. The sensitivity graphs are shown in Figure 4. For the "Activation" dataset:

- 938 • **gradient\_std\_multiplier:** Increasing this parameter until a certain range had a stable effect  
 939 on performance. Beyond this point, AUPRC for warm and drug cold starts began to drop,  
 940 while protein cold start AUPRC remained stable. Simultaneously, AUROC for drug and  
 941 warm starts decreased, while protein cold start AUROC increased.
- 942 • **gradient\_history\_length:** Model performance was relatively stable with respect to its value  
 943 increase across all evaluation setups.

## 944 C.5 VOLUME COMPUTATION DETAILS

945 The GRAM4Modal and GRAM3Modal functions compute volumes using Gram matrix determinants.  
 946 For anchor features  $f^a$  and target features  $\{f^{t_1}, f^{t_2}, f^{t_3}\}$ , the  $4 \times 4$  Gram matrix  $G$  has entries  
 947  $G_{kj} = \langle f^k, f^j \rangle$ . The volume is computed as  $V = \sqrt{|\det(G)|}$ , then converted to similarity via  
 948 negative volume scaling:  $S = -V/\tau$ .

949 Algorithm 2 implements our gradient-informed adaptive modality selection strategy, which maintains  
 950 consistency between forward  $\mathcal{L}_{\text{vol}}^{\rightarrow}$  and reverse  $\mathcal{L}_{\text{vol}}^{\leftarrow}$  contrastive computations by using a single  
 951 drop decision per forward pass.

## 952 C.6 NEGATIVE SAMPLING STRATEGIES

953 We construct negative samples by fixing all but one modality, producing hard negatives in which  
 954 only a single modality is mismatched while the remaining modalities are aligned. This single-

972 modality perturbation yields a more challenging learning signal, as the model must distinguish the  
 973 fully aligned positive case from near-aligned negatives. To assess alternative strategies, we also  
 974 evaluate an aggressive multi-domain negative-sampling scheme in which multiple modalities are  
 975 perturbed simultaneously. Specifically, for each sample  $i$  in batch  $B$ , we generate negative samples  
 976 by permuting all modalities. Results on the Activation dataset comparing cross-negative sampling to  
 977 the current strategy under the volume-loss pretraining setting are reported in Table 9 and Table 10.

980  
 981 **Table 9: Model performance comparison on activation dataset in terms of AUROC when using  
 different negative sample strategies: Cross-Negative vs Current-Negative (Mean±Std).**

982 Split Type	983 Cross-Negative	984 Current-Negative
983 Warm start	984 $0.9142 \pm 0.0071$	985 $0.9142 \pm 0.0078$
984 Drug cold start	985 $0.9164 \pm 0.0093$	986 $0.9125 \pm 0.0068$
985 Target cold start	986 $0.8388 \pm 0.0272$	987 $0.8335 \pm 0.0258$

988  
 989 **Table 10: Model performance comparison on activation dataset in terms of AUPR when using dif-  
 990 ferent negative sample strategies: Cross-Negative vs Current-Negative (Mean±Std).**

992 Split Type	993 Cross-Negative	994 Current-Negative
993 Warm start	994 $0.6326 \pm 0.0232$	995 $0.6424 \pm 0.0221$
994 Drug cold start	995 $0.6239 \pm 0.0245$	996 $0.6278 \pm 0.0222$
995 Target cold start	996 $0.4618 \pm 0.0313$	997 $0.4497 \pm 0.0374$

### 1000 C.7 SENSITIVITY TO BATCH SIZE AND IN-BATCH NEGATIVES

1001 Our method uses the standard in-batch negative formulation: for a batch of size  $N$ , each sample  
 1002 uses the other  $N - 1$  samples as negatives. We do not use stabilization techniques such as memory  
 1003 banks. To test sensitivity, we varied the per-GPU batch size from 32 up to 512. The main paper  
 1004 results used a batch size of 1280. As expected, performance generally improves with larger batch  
 1005 sizes, since more in-batch negatives benefit the contrastive and volume losses. The corresponding  
 1006 results are shown in table 11 and 12.

1008 **Table 11: AUROC Performance vs. Batch Size (Activation Dataset)**

1010 Batch Size	1011 Warm Start	1012 Drug Cold Start	1013 Target Cold Start
1011 32	1012 0.901	1013 0.906	1014 0.819
1012 128	1013 0.905	1014 0.906	1015 0.821
1013 512	1014 0.916	1015 0.918	1016 0.839
1014 1280	1015 0.914	1016 0.913	1017 0.834

1018 Note that the 512 batch size results are very close to the 1280 results, suggesting performance may  
 1019 begin to saturate beyond batch size 512.

### 1021 C.8 DOWNSTREAM TASK ARCHITECTURE

1022 For drug-target interaction (DTI) prediction evaluation, we employ a lightweight classification archi-  
 1023 tecture that leverages the pre-trained embeddings from our four-modal framework. The downstream  
 1024 architecture is detailed in Algorithm 3 and uses only the drug (SMILES) and protein modalities  
 1025 relevant for binding prediction.

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

Table 12: AUPRC Performance vs. Batch Size (Activation Dataset)

Batch Size	Warm Start	Drug Cold Start	Target Cold Start
32	0.615	0.607	0.432
128	0.625	0.619	0.438
512	0.642	0.629	0.458
1280	0.642	0.628	0.450

**Algorithm 3** Drug-Target Interaction Prediction

**Require:** Pre-trained embeddings  $f^s, f^p \in \mathbb{R}^{512}$

**Require:** Drug-protein pair  $(x_i^s, x_j^p)$ , binding label  $y_{ij} \in \{0, 1\}$

**Ensure:** Binding prediction  $\hat{y}_{ij}$

- 1:  $f_i^s \leftarrow \text{FROZEN}(F_\phi^s(E_s(x_i^s)))$  {Use pre-trained SMILES embedding}
- 2:  $f_j^p \leftarrow \text{FROZEN}(F_\phi^p(E_p(x_j^p)))$  {Use pre-trained protein embedding}
- 3:  $f^{\text{concat}} \leftarrow [f_i^s; f_j^p] \in \mathbb{R}^{1024}$  {Concatenate embeddings}
- 4:  $h_1 \leftarrow \text{ReLU}(\text{Linear}_{1024 \rightarrow 512}(f^{\text{concat}}))$
- 5:  $h_1 \leftarrow \text{Dropout}_{0.3}(h_1)$
- 6:  $h_2 \leftarrow \text{ReLU}(\text{Linear}_{512 \rightarrow 256}(h_1))$
- 7:  $h_2 \leftarrow \text{Dropout}_{0.3}(h_2)$
- 8:  $\text{logits} \leftarrow \text{Linear}_{256 \rightarrow 2}(h_2)$
- 9:  $\hat{y}_{ij} \leftarrow \arg \max(\text{softmax}(\text{logits}))$
- 10: **return**  $\hat{y}_{ij}$

**C.9 EVALUATION METRICS**

We employ five standard binary classification metrics to comprehensively assess DTI prediction performance. Given the confusion matrix with true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN), the metrics are defined as follows:

**Area Under ROC Curve (AUROC)** AUROC measures the model’s ability to discriminate between positive and negative classes across all classification thresholds:

$$\text{AUROC} = \int_0^1 \text{TPR}(\text{FPR}^{-1}(t)) dt \quad (9)$$

where  $\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}$  and  $\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}}$ .

**Area Under Precision-Recall Curve (AUPRC)** AUPRC is particularly informative for imbalanced datasets and measures performance across different precision-recall trade-offs:

$$\text{AUPRC} = \int_0^1 \text{Precision}(\text{Recall}^{-1}(t)) dt \quad (10)$$

where  $\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$  and  $\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$ .

**Sensitivity (Recall)** Sensitivity measures the proportion of actual positive cases correctly identified:

$$\text{Sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (11)$$

**F1-Score** F1-score provides the harmonic mean of precision and recall, balancing both measures:

$$\text{F1} = 2 \cdot \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{2 \cdot \text{TP}}{2 \cdot \text{TP} + \text{FP} + \text{FN}} \quad (12)$$

**Accuracy** Accuracy measures the overall proportion of correct predictions:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (13)$$

## 1080 D ABLATION STUDY

1081  
 1082 To complement the ablation study presented in Section 4.5 on the Activation dataset, we provide  
 1083 additional comprehensive ablation experiments on the Yamanishi 08 dataset in Figure ???. This  
 1084 additional evaluation allows us to assess the generalizability of our component contributions across  
 1085 different datasets and task characteristics.

### 1087 D.1 EXPERIMENTAL SETUP

1088 The ablation study on Yamanishi 08 follows the same experimental configuration as described in  
 1089 Section 4.5, evaluating five distinct setups:

- 1091 • **Exp 1:** Full GRAM-DTI model with all components and adaptive modality dropout
- 1092 • **Exp 2:** Training without Gramian volume-based loss ( $L = \lambda_2 L_{\text{bi}} + \lambda_3 L_{\text{IC50}}$ )
- 1093 • **Exp 3:** Training without bimodal contrastive loss ( $L = \lambda_1 L_{\text{vol}} + \lambda_3 L_{\text{IC50}}$ )
- 1094 • **Exp 4:** Training without IC50 auxiliary supervision ( $L = \lambda_1 L_{\text{vol}} + \lambda_2 L_{\text{bi}}$ )
- 1095 • **Exp 5:** Training with full objective but without adaptive modality dropout

### 1098 D.2 RESULTS ANALYSIS

1100 The results on Yamanishi 08, shown in Figure 5, demonstrate consistent patterns with those observed  
 1101 on the Activation dataset, confirming the robustness of our design choices across different datasets.

1102 **Consistent Superior Performance of Full Model:** Across all three data splitting scenarios (warm  
 1103 start, drug cold start, target cold start) and five evaluation metrics (AUROC, AUPRC, Sensitivity,  
 1104 F1, Accuracy), the full GRAM-DTI model (Exp 1) generally achieves the highest performance,  
 1105 demonstrating the synergistic benefit of all proposed components.

- 1107 1. The volume-based multimodal alignment provides substantial benefits over traditional pair-  
 1108 wise approaches
- 1109 2. Adaptive modality dropout prevents overfitting and improves generalization
- 1110 3. IC50 auxiliary supervision enhances biological relevance of learned representations
- 1111 4. The synergistic combination of all components yields optimal performance

1113 These consistent findings across different datasets and evaluation scenarios validate the generaliz-  
 1114 ability of our GRAM-DTI framework design principles.

## 1116 E ADDITIONAL EXPERIMENTAL DETAILS

1120 Table 13: Performance metrics with standard deviations for GRAM-DTI across all evaluation  
 1121 datasets and data splitting scenarios. Results are reported as mean  $\pm$  standard deviation across cross-  
 1122 validation folds.

1123 <b>Dataset</b>	1124 <b>Split Type</b>	1125 <b>AUROC <math>\uparrow</math></b>	1126 <b>AUPRC <math>\uparrow</math></b>	1127 <b>Sensitivity <math>\uparrow</math></b>	1128 <b>F1 <math>\uparrow</math></b>	1129 <b>Accuracy <math>\uparrow</math></b>
1130 Yamanishi_08	warm start	0.9771 $\pm$ 0.0042	0.9036 $\pm$ 0.0079	0.7954 $\pm$ 0.0152	0.8353 $\pm$ 0.0096	0.9715 $\pm$ 0.0015
	drug cold start	0.8279 $\pm$ 0.0285	0.4404 $\pm$ 0.0662	0.2020 $\pm$ 0.0575	0.3090 $\pm$ 0.0693	0.9193 $\pm$ 0.0134
	target cold start	0.9553 $\pm$ 0.0155	0.8494 $\pm$ 0.0312	0.7189 $\pm$ 0.0453	0.7840 $\pm$ 0.0285	0.9643 $\pm$ 0.0042
1131 Hetionet	warm start	0.9808 $\pm$ 0.0011	0.8586 $\pm$ 0.0082	0.7580 $\pm$ 0.0085	0.7891 $\pm$ 0.0065	0.9632 $\pm$ 0.0010
	drug cold start	0.8550 $\pm$ 0.0385	0.5291 $\pm$ 0.0626	0.2981 $\pm$ 0.0645	0.4227 $\pm$ 0.0619	0.9273 $\pm$ 0.0131
	target cold start	0.9210 $\pm$ 0.0079	0.6258 $\pm$ 0.0239	0.4569 $\pm$ 0.0448	0.5502 $\pm$ 0.0319	0.9325 $\pm$ 0.0038
1132 Activation	warm start	0.9142 $\pm$ 0.0078	0.6424 $\pm$ 0.0221	0.5155 $\pm$ 0.0240	0.5950 $\pm$ 0.0075	0.9364 $\pm$ 0.0026
	drug cold start	0.9125 $\pm$ 0.0068	0.6278 $\pm$ 0.0222	0.5135 $\pm$ 0.0349	0.5879 $\pm$ 0.0186	0.9347 $\pm$ 0.0030
	target cold start	0.8335 $\pm$ 0.0258	0.4497 $\pm$ 0.0374	0.2451 $\pm$ 0.0591	0.3447 $\pm$ 0.0620	0.9168 $\pm$ 0.0104
1133 Inhibition	warm start	0.9491 $\pm$ 0.0018	0.7849 $\pm$ 0.0061	0.6588 $\pm$ 0.0109	0.7202 $\pm$ 0.0061	0.9535 $\pm$ 0.0013
	drug cold start	0.9398 $\pm$ 0.0018	0.7555 $\pm$ 0.0034	0.5949 $\pm$ 0.0176	0.6801 $\pm$ 0.0081	0.9492 $\pm$ 0.0011
	target cold start	0.8234 $\pm$ 0.0218	0.4641 $\pm$ 0.0559	0.2584 $\pm$ 0.0827	0.3687 $\pm$ 0.0872	0.9220 $\pm$ 0.0087

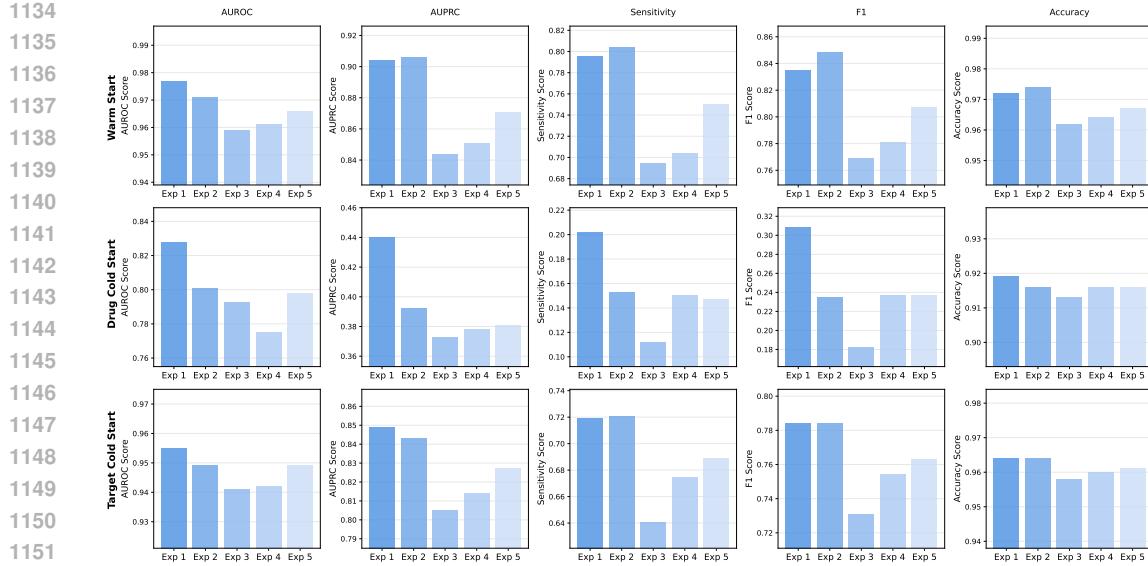


Figure 5: Ablation study results on the Yamanishi 08 dataset across five experimental configurations and three data splitting scenarios. The full GRAM-DTI model (Exp 1) consistently outperforms variants with removed components across most metrics and scenarios, demonstrating the robust contribution of each training objective component. Results complement those shown in Figure 2 (Activation dataset) and confirm the generalizability of our design choices across different DTI prediction benchmarks.

## E.1 STANDARD DEVIATION RESULTS FOR MAIN PERFORMANCE COMPARISON

Table 13 provides comprehensive performance statistics for GRAM-DTI, including standard deviations across all evaluation metrics, datasets, and data splitting scenarios. These detailed statistics demonstrate the stability and reliability of our approach across cross-validation folds.

## E.2 PERFORMANCE ANALYSIS WHEN CERTAIN MODALITIES ARE MISSING AT PRETRAINING TIME

Bringing in as many relevant modalities in the pretraining would help learn better representation for the corresponding downstream task. However, what if during training, not all the modalities are available? We investigated this question by considering scenario where certain modality is not available during pretraining. For drug-target interaction tasks, the drug and protein modalities are crucial. If either is unavailable during pretraining, the model cannot learn joint embeddings for interacting pairs. If for one the missing modality, what one can do is use only the embedding obtained from original encoders (ESM2/molformer) without further fine tuning it with the contrastive learning, this scenario will effectively falls back to a regime similar to DTIAM, where drug and protein embeddings are learned separately. This highlights a key strength of GRAM-DTI: by pretraining with both modalities present, it learns aligned embeddings that improve downstream performance. For auxiliary modalities such as functional descriptors and HTA:

- If both are missing during pretraining, this is equivalent to the ablation study (Exp2), where the volume-loss component is removed (as those two modality anticipate in the training through volume loss). In this case, we observe a slight decrease in performance, indicating that these modalities provide useful signals for alignment.
- If only one is missing, as shown in Table 14 and 15 the model still benefits from the available modality, with a moderate drop in performance. This demonstrates that GRAM-DTI

1188  
 1189 can gracefully handle partial modality availability, but full multi-modal pretraining yields  
 1190 the strongest embeddings.  
 1191

1192 Table 14: Model performance on the Activation dataset in terms of *AUROC* for different modality  
 1193 configurations: 3-mod-no-text: all modality except functional descriptors are available during pre-  
 1194 training, 3-mod-no-HTA: all modality except HTA are available during pretraining , 4-mod: all four  
 1195 modalities are available during pretraining (current setup).

Split Type	3-mod-no-text	3-mod-no-HTA	4-mod
Warm start	0.907	0.901	0.914
Drug cold start	0.907	0.903	0.913
Target cold start	0.828	0.821	0.834

1200  
 1201  
 1202 Table 15: Model performance on the Activation dataset in terms of *AUPRC* for different modality  
 1203 configurations: 3-mod-no-text: all modality except functional descriptors are available during pre-  
 1204 training, 3-mod-no-HTA: all modality except HTA are available during pretraining , 4-mod: all four  
 1205 modalities are available during pretraining (current setup).

Split Type	3-mod-no-text	3-mod-no-HTA	4-mod
Warm start	0.609	0.606	0.642
Drug cold start	0.615	0.611	0.628
Target cold start	0.437	0.440	0.450

### 1214 E.3 OVERLAP OF ENTITIES ANALYSIS BETWEEN PRETRAINING AND DOWNSTREAM TASK

1215 To verify that our method does not memorize entity-specific patterns, we conducted an overlap  
 1216 analysis on the Activation dataset between pretraining and downstream task dataset (Other datasets  
 1217 are shown in Table 17). The results revealed 236 overlapping proteins and 314 overlapping SMILES.  
 1218 We removed all pairs containing these overlapping entities from the pretraining data, resulting in  
 1219 6,065 exact (SMILES, protein) pairs removed (11.9% of pretraining data). The results are shown in  
 1220 Table 16. Despite removing nearly 12% of the pretraining data, the performance drops are modest  
 1221

1223 Table 16: Performance before and after cleaning on the Activation dataset.

Split Type	Metric	Before Cleaning	After Cleaning	$\Delta$
Warm Start	<i>AUROC</i>	0.914	0.901	-0.013
	<i>AUPRC</i>	0.642	0.613	-0.029
Drug Cold Start	<i>AUROC</i>	0.913	0.905	-0.008
	<i>AUPRC</i>	0.628	0.624	-0.004
Target Cold Start	<i>AUROC</i>	0.834	0.795	-0.039
	<i>AUPRC</i>	0.450	0.389	-0.061

1232 across all splits (0.01-0.03). This demonstrates that our model’s strong performance is not primarily  
 1233 driven by memorizing exact pairs, and further validates our cold-start claims:  
 1234

- 1235 • Entity overlap contributes to performance but is not the dominant factor
- 1236 • The 4-modal learning framework captures transferable molecular representations rather  
 1237 than memorizing specific entity combinations
- 1238 • Drug cold-start generalization is particularly robust ( $\Delta$  AUROC = -0.008), showing mini-  
 1239 mal sensitivity to entity overlap

1242

1243

Table 17: Overlap analysis between pretraining and downstream datasets.

1244

1245

1246

1247

1248

1249

1250

1251

## E.4 MODEL PERFORMANCE SENSITIVITY TO THE NEGATIVE-SAMPLE RATIO IN THE DOWNSTREAM TASK

1252

1253

All baselines generate negative samples at a 1:10 ratio relative to positive samples across datasets; to ensure a fair comparison, we adopt the same setup. To evaluate sensitivity to this choice, we also report results using alternative negative-sample ratios. In table 18, we show our model performance as well as best baseline (DTIAM) on the Activation dataset when negative samples are generated at various ratio with respect to positive samples.

1254

1255

1256

1257

1258

1259

1260

1261

Table 18: Performance Comparison of GRAM-DTI and Baseline on Activation Dataset under Different Ratios

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

## E.5 HANDLING PARTIAL-MODALITY DATA DURING PRETRAINING

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

we can extend pretraining to include samples with missing modalities, which would substantially increase the size of our training set. To assess whether all modalities are beneficial, our current pretraining phase includes only samples in which all four modalities are present, a choice that significantly limits the dataset. As a proof of concept, we evaluated whether including samples with only a subset of modalities improves downstream performance. From the pretraining dataset we created:

- Fully Observed (80%): 80% of the original data, kept unchanged.
- Partially Observed (20%): the remaining 20% where we randomly dropped one modality.

We compared training on only the 80% fully-observed subset vs. training on the full 100% dataset (80% full + 20% partial) using a masked-volume loss for the partial samples. This simulates the

Dataset	Total Pairs	Overlapping	Percentage	Proteins	SMILES
activation	50,968	6,065	11.90%	236	314
Hetionet	50,968	42,242	82.88%	1,936	853
inhibition	50,968	36,372	71.36%	860	1,382
yamanishi_08	50,968	20,223	39.68%	556	344

1296 setup suggested by the reviewer and shows how pretraining can be expanded when some modalities  
 1297 are missing at random.  
 1298

1299  
1300 **Table 19: Performance with Partial-Modality (“Masked-Volume”) Training**

1301 Split Type	1302 Metric	1303 Fully Observed Only (80% data)	1304 Full + Partial (100% data)
1302 Warm Start	AUROC	0.905	0.912
	AUPRC	0.627	0.634
1304 Drug Cold Start	AUROC	0.903	0.907
	AUPRC	0.613	0.615
1306 Target Cold Start	AUROC	0.791	0.828
	AUPRC	0.422	0.437

1309  
1310 As the table 19 shows, incorporating the 20% partial data via masked-volume training improves  
 1311 performance across all metrics and splits, with a notable improvement on the Target Cold Start  
 1312 (AUROC 0.828 vs. 0.791).  
 1313  
 1314

## 1315 E.6 ANALYSIS OF DROPOUT VS. WEIGHTING STRATEGIES

1316 To validate our Gradient-Informed Modality Dropout strategy, we compared it against two alterna-  
 1317 tive “soft” balancing mechanisms on the Activation dataset:  
 1318

- 1319 • **Weighted-Modality Gradients:** Instead of dropping a modality, we scale its gradient by  
 1320 the inverse of its norm with probability  $p_{drop}$ .  
 1321
- 1322 • **Standard Weighted Loss:** We assign learnable weights to each modality’s loss term to  
 1323 balance contributions without dropout.  
 1324

1325 As shown in Table 20, our probabilistic dropout strategy achieves the best performance. We hypo-  
 1326 thesize that probabilistically removing modalities forces the model to find alternative distinct paths for  
 1327 reasoning in the joint embedding space, acting as a stronger regularizer than soft weighting.  
 1328

1329 **Table 20: Comparison of Modality Balancing Strategies on the Activation Dataset. Our hard dropout**  
 1330 **strategy outperforms soft weighting approaches.**

1331 Strategy	1332 Split Type	1333 AUROC	1334 AUPRC
1332 Gradient-Informed Dropout (Ours)	1333 Warm Start	<b>0.914</b>	<b>0.642</b>
	1333 Drug Cold Start	<b>0.913</b>	<b>0.628</b>
	1333 Target Cold Start	<b>0.834</b>	<b>0.450</b>
1335 Weighted Gradients	1335 Warm Start	0.909	0.618
	1335 Drug Cold Start	0.910	0.624
	1335 Protein Cold Start	0.828	0.445
1338 Standard Weighted Loss	1338 Warm Start	0.901	0.621
	1338 Drug Cold Start	0.892	0.619
	1338 Target Cold Start	0.814	0.440

## 1343 E.7 EXPERIMENTS WITH STRONGER MOLECULAR ENCODERS

1344 Our default GRAM-DTI implementation uses MolFormer for computational efficiency. To demon-  
 1345 strate the framework’s extensibility, we replaced MolFormer with two larger, more advanced en-  
 1346 coders: Uni-Mol2 (84M parameters) and BioT5+. As shown in Table 21, utilizing stronger encoders  
 1347 consistently improves performance across all splits, particularly in the challenging target Cold start  
 1348 scenario. It shows that stronger encoders yield consistently better results.  
 1349

1350

1351 Table 21: Sensitivity analysis using advanced molecular encoders on the Activation dataset

Encoder	Split Type	AUROC	AUPRC
MolFormer (Original)	Warm Start	0.9142	0.6424
	Drug Cold Start	0.9125	0.6278
	Target Cold Start	0.8335	0.4497
Uni-Mol2 (84M)	Warm Start	<b>0.9280</b>	0.6768
	Drug Cold Start	<b>0.9270</b>	0.6658
	Target Cold Start	<b>0.8642</b>	<b>0.4848</b>
BioT5+	Warm Start	0.9273	<b>0.6828</b>
	Drug Cold Start	0.9254	<b>0.6840</b>
	Target Cold Start	0.8577	0.4805

1363

1364

## 1365 E.8 FALSE NEGATIVE CASE ANALYSIS

1366

1367 To understand where the model fails, we systematically identified the top “hardest” false negatives  
 1368 in the Activation dataset—pairs where the model predicted a strong negative signal despite a positive  
 1369 ground truth label. These are listed in table 22 below:

1370

1371

1372

Table 22: Top-10 False Negative Pairs (Drug ID &amp; Target ID)

Rank	Drug ID	Target ID
1	D0NY1R	T36075
2	D08FKH	T12475
3	D0G2VT	T59604
4	D0JB3H	T88505
5	D0L5WA	T28893
6	D0K8NR	T72458
7	D03LQC	T52522
8	D03XIS	T92076
9	D07QAK	T28893
10	D0AJ2T	T88505

1384

1385

1386

## 1387 E.9 MoA TASK ADDITIONAL BASELINES

1388

1389 To extend the set of baselines beyond those used in the DTIAM study for the MoA task, we included  
 1390 two additional methods: DeepDTA (Öztürk et al., 2018) and GraphDTA (Nguyen et al., 2021).  
 1391 Although DTIAM remains the strongest baseline overall, our model GRAM-DTI achieves superior  
 1392 performance in most evaluation settings. The results are shown in the table 23.

1393

1394

## E.10 SIGNIFICANCE TEST

1395

1396

1397 We ran one-sided Welch t-tests to assess whether GRAM-DTI outperforms DTIAM (the strongest  
 1398 baseline). Tests were computed from summary statistics (means and standard deviations) using n=10  
 1399 (10 folds) for Yamanishi\_08 and Hetionet dataset and n=5 (5 folds) for Activation and Inhibition  
 1400 dataset with the one-sided hypothesis H1: GRAM-DTI > DTIAM. Table 24 and Table 25 report  
 1401 the corresponding p values for the MoA and DTI tasks, respectively, and Tables 26 and 27 show  
 1402 the zero-shot retrieval results. Cells highlighted in light blue indicate the better method in each  
 1403 row. Note that we performed a total of 48 tests. To control for multiple comparisons we applied a  
 1404 Bonferroni correction and used an adjusted significance threshold of:  $p_{adjusted} = \frac{0.05}{48} \approx 0.00104$ ,  
 1405 rather than the conventional  $p = 0.05$ . p-values highlighted in light blue indicate  $p < 0.00104$ .

1404  
1405 Table 23: Performance comparison between GRAM-DTI and state-of-the-art baselines (DeepDTA,  
1406 GraphDTA, AI-DTI, DTIAM) on MoA prediction tasks across multiple datasets and data splitting  
1407 scenario. GRAM-DTI demonstrates superior performance in most evaluation settings.  $\dagger$  indicates  
1408 reproduced results; other results are from baseline papers. **Bold** denotes best performance.

Data	Metric	Scenario	DeepDTA $\dagger$	GraphDTA $\dagger$	AI-DTI	DTIAM $\dagger$	GRAM-DTI
Activation	AUPR	Warm Start	0.246 $\pm$ 0.0232	0.282 $\pm$ 0.0240	0.583	0.623 $\pm$ 0.0245	<b>0.642<math>\pm</math>0.0221</b>
		Drug Cold Start	0.255 $\pm$ 0.0209	0.298 $\pm$ 0.0195	0.550	0.611 $\pm$ 0.0252	<b>0.628<math>\pm</math>0.0222</b>
		Cold Start	0.109 $\pm$ 0.0163	0.124 $\pm$ 0.0175	0.219	0.391 $\pm$ 0.0320	<b>0.450<math>\pm</math>0.0374</b>
	AUROC	Warm Start	0.759 $\pm$ 0.0200	0.784 $\pm$ 0.0185	0.888	0.903 $\pm$ 0.0088	<b>0.914<math>\pm</math>0.0078</b>
		Drug Cold Start	0.765 $\pm$ 0.0059	0.796 $\pm$ 0.0062	0.879	0.907 $\pm$ 0.0076	<b>0.913<math>\pm</math>0.0068</b>
		Cold Start	0.573 $\pm$ 0.0241	0.588 $\pm$ 0.0255	0.652	0.792 $\pm$ 0.0240	<b>0.834<math>\pm</math>0.0258</b>
Inhibition	AUPR	Warm Start	0.542 $\pm$ 0.0195	0.585 $\pm$ 0.0280	0.840	<b>0.845<math>\pm</math>0.0070</b>	0.785 $\pm$ 0.0061
		Drug Cold Start	0.531 $\pm$ 0.0170	0.592 $\pm$ 0.0195	<b>0.830</b>	0.731 $\pm$ 0.0045	0.756 $\pm$ 0.0034
		Cold Start	0.265 $\pm$ 0.0210	0.284 $\pm$ 0.0312	0.215	0.445 $\pm$ 0.0620	<b>0.464<math>\pm</math>0.0559</b>
	AUROC	Warm Start	0.854 $\pm$ 0.0105	0.872 $\pm$ 0.0098	0.952	<b>0.954<math>\pm</math>0.0025</b>	0.949 $\pm$ 0.0018
		Drug Cold Start	0.849 $\pm$ 0.0185	0.876 $\pm$ 0.0115	<b>0.948</b>	0.921 $\pm$ 0.0028	0.940 $\pm$ 0.0018
		Cold Start	0.635 $\pm$ 0.0220	0.649 $\pm$ 0.0117	0.605	0.819 $\pm$ 0.0205	<b>0.823<math>\pm</math>0.0028</b>

1421  
1422 Table 24: Significance test on the MoA task: means ( $\pm$  std) and one-sided  $p$ -values for  $H_1$ : GRAM  
1423 > DTIAM (independent Welch test). One-sided  $p$ -values with  $p < 0.00104$  are highlighted in light  
1424 blue.

Dataset	Metric	Scenario	DTIAM ( $\mu \pm \sigma$ )	GRAM ( $\mu \pm \sigma$ )	one-sided $p$ value
Activation	AUPR	Warm start	0.623 $\pm$ 0.0245	0.642 $\pm$ 0.0221	0.1171
		Drug cold start	0.611 $\pm$ 0.0252	0.628 $\pm$ 0.0222	0.1455
		Target cold start	0.391 $\pm$ 0.0320	0.450 $\pm$ 0.0374	0.0143
	AUROC	Warm start	0.903 $\pm$ 0.0088	0.914 $\pm$ 0.0078	0.0352
		Drug cold start	0.907 $\pm$ 0.0076	0.913 $\pm$ 0.0068	0.1126
		Target cold start	0.792 $\pm$ 0.0240	0.834 $\pm$ 0.0258	0.0144
Inhibition	AUPR	Warm start	0.845 $\pm$ 0.0070	0.785 $\pm$ 0.0061	1.0000
		Drug cold start	0.731 $\pm$ 0.0045	0.756 $\pm$ 0.0034	< 0.0001
		Target cold start	0.445 $\pm$ 0.0620	0.464 $\pm$ 0.0559	0.3123
	AUROC	Warm start	0.954 $\pm$ 0.0025	0.949 $\pm$ 0.0018	0.9961
		Drug cold start	0.921 $\pm$ 0.0028	0.940 $\pm$ 0.0018	< 0.0001
		Target cold start	0.819 $\pm$ 0.0205	0.823 $\pm$ 0.0028	0.3435

## 1440 E.11 ZERO-SHOT RETRIEVAL TASK METHODOLOGY

1441 This section provides detailed methodology for the zero-shot retrieval experiments presented in  
1442 Section 4.3 of the main text.

### 1445 E.11.1 TASK FORMULATION

1446 The zero-shot retrieval task evaluates GRAM-DTI’s ability to identify relevant drug-target pairs using  
1447 only the learned multimodal representations, without any task-specific fine-tuning. We formulate  
1448 two complementary retrieval scenarios:

- 1449 • **Drug-to-Protein Retrieval (S $\rightarrow$ P):** Given a query drug (SMILES representation), retrieve  
1450 the most relevant target proteins from a candidate set.
- 1451 • **Protein-to-Drug Retrieval (P $\rightarrow$ S):** Given a query protein (sequence representation), re-  
1452 trieve the most relevant drugs from a candidate set.

### 1453 E.11.2 EXPERIMENTAL SETUP

1454 For each dataset, we construct retrieval queries and candidate pools as follows:

1458  
1459  
1460  
1461Table 25: Significance test on the DTA task: means ( $\pm$  std) and one-sided  $p$ -values for  $H_1$ : GRAM > DTIAM (independent Welch test). One-sided  $p$ -values with  $p < 0.00104$  are highlighted in light blue.

Dataset	Metric	Scenario	DTIAM ( $\mu \pm \sigma$ )	GRAM ( $\mu \pm \sigma$ )	one-sided $p$ value
Yamanishi_08	AUPR	Warm start	$0.901 \pm 0.0085$	$0.904 \pm 0.0079$	0.2122
		Drug cold start	$0.439 \pm 0.0580$	$0.440 \pm 0.0662$	0.4859
		Target cold start	$0.844 \pm 0.0350$	$0.849 \pm 0.0312$	0.3700
	AUROC	Warm start	$0.967 \pm 0.0050$	$0.977 \pm 0.0042$	< 0.0001
		Drug cold start	$0.818 \pm 0.0255$	$0.828 \pm 0.0285$	0.2096
		Target cold start	$0.941 \pm 0.0180$	$0.955 \pm 0.0155$	0.0395
Hetionet	AUPR	Warm start	$0.879 \pm 0.0095$	$0.859 \pm 0.0082$	0.9984
		Drug cold start	$0.514 \pm 0.0680$	$0.529 \pm 0.0626$	0.3070
		Target cold start	$0.625 \pm 0.0210$	$0.626 \pm 0.0239$	0.4610
	AUROC	Warm start	$0.957 \pm 0.0015$	$0.981 \pm 0.0011$	< 0.0001
		Drug cold start	$0.752 \pm 0.0355$	$0.855 \pm 0.0385$	< 0.0001
		Target cold start	$0.917 \pm 0.0090$	$0.921 \pm 0.0079$	0.1525

1470  
1471  
1472  
1473  
1474  
1475  
1476Table 26: Significance test on Zero-shot retrieval task (Yamanishi\_08 and Hetionet): one-sided Welch  $t$ -test ( $H_1$ : GRAM > DTIAM), values  $p < 0.00104$  are highlighted.

Direction	Metric	Yamanishi_08			Hetionet		
		DTIAM	GRAM	$p$	DTIAM	GRAM	$p$
S→P	R@1	$0.0038 \pm 0.0004$	<b><math>0.0465 \pm 0.0027</math></b>	< 0.0001	$0.0043 \pm 0.0002$	<b><math>0.0331 \pm 0.0038</math></b>	< 0.0001
	R@10	$0.0341 \pm 0.0042$	<b><math>0.1691 \pm 0.0084</math></b>	< 0.0001	$0.0434 \pm 0.0051$	<b><math>0.1340 \pm 0.0025</math></b>	< 0.0001
	R@100	$0.1960 \pm 0.0181$	<b><math>0.4449 \pm 0.0075</math></b>	< 0.0001	$0.2066 \pm 0.0109$	<b><math>0.3616 \pm 0.0063</math></b>	< 0.0001
P→S	R@1	$0.0040 \pm 0.0002$	<b><math>0.0742 \pm 0.0120</math></b>	< 0.0001	$0.0404 \pm 0.0028$	$0.0236 \pm 0.0010$	1.0000
	R@10	$0.0849 \pm 0.0089$	<b><math>0.2465 \pm 0.0256</math></b>	< 0.0001	$0.1319 \pm 0.0095$	$0.1049 \pm 0.0055$	1.0000
	R@100	$0.3670 \pm 0.0186$	<b><math>0.5540 \pm 0.0148</math></b>	< 0.0001	$0.3632 \pm 0.0474$	$0.3841 \pm 0.0082$	0.1900

1487  
1488Table 27: Significance test on Zero-shot retrieval task (Activation and Inhibition): One-sided Welch  $t$ -test ( $H_1$ : GRAM > DTIAM),  $p < 0.00104$  values are highlighted.

Direction	Metric	Activation			Inhibition		
		DTIAM	GRAM	$p$	DTIAM	GRAM	$p$
S→P	R@1	$0.0028 \pm 0.0002$	<b><math>0.0136 \pm 0.0011</math></b>	< 0.0001	$0.0004 \pm 0.0000$	<b><math>0.0055 \pm 0.0003</math></b>	< 0.0001
	R@10	$0.0266 \pm 0.0037$	<b><math>0.1020 \pm 0.0067</math></b>	< 0.0001	$0.0097 \pm 0.0006$	<b><math>0.0337 \pm 0.0011</math></b>	< 0.0001
	R@100	$0.3184 \pm 0.0229$	<b><math>0.5688 \pm 0.0172</math></b>	< 0.0001	$0.1036 \pm 0.0104$	<b><math>0.1994 \pm 0.0018</math></b>	< 0.0001
P→S	R@1	$0.0071 \pm 0.0008$	<b><math>0.0370 \pm 0.0069</math></b>	< 0.0001	$0.0000 \pm 0.0000$	<b><math>0.0221 \pm 0.0061</math></b>	< 0.0001
	R@10	$0.0463 \pm 0.0050$	<b><math>0.2454 \pm 0.0142</math></b>	< 0.0001	$0.0028 \pm 0.0004$	<b><math>0.0819 \pm 0.0065</math></b>	< 0.0001
	R@100	$0.2206 \pm 0.0264$	<b><math>0.6029 \pm 0.0231</math></b>	< 0.0001	$0.0588 \pm 0.0049$	<b><math>0.2325 \pm 0.0094</math></b>	< 0.0001

1496

**Query and Candidate Construction:**1497  
1498  
1499  
1500  
1501  
1502  
1503

- For each known drug-target interaction ( $d_i, p_j$ ) in the data set, we treat  $d_i$  as a query and all proteins in the dataset as candidates for S→P retrieval
- Similarly, we treat  $p_j$  as a query and all drugs as candidates for P→S retrieval
- Ground truth relevance is determined by known interactions in the original datasets

1504

**Embedding Generation:** We generate embeddings using the pre-trained GRAM-DTI framework:1505  
1506  
1507  
1508  
1509  
1510  
1511

- SMILES sequences are encoded using MoLFormer-XL, producing 768-dimensional representations
- Protein sequences are encoded using ESM-2, producing 1280-dimensional representations
- Both modalities are projected to a shared 512-dimensional space using trained projectors from the multimodal pre-training phase
- All embeddings are L2-normalized for cosine similarity computation

1512  
1513 **Similarity Computation:** We compute cosine similarity between query and candidate representations using the projected embeddings:  
1514

$$\text{sim}(q, c) = \frac{f_q \cdot f_c}{\|f_q\| \|f_c\|} \quad (14)$$

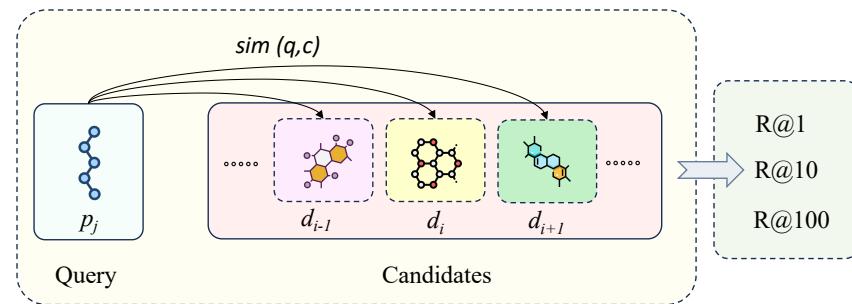
1518 where  $f_q$  and  $f_c$  are the normalized projected embeddings for query  $q$  and candidate  $c$ , respectively.  
1519

#### 1520 Ranking and Evaluation:

- 1521 1. For each query, we rank all candidates by their similarity scores in descending order
- 1522 2. We evaluate retrieval performance using standard ranking metrics:
  - 1524 • **Recall@1 (R@1):** Proportion of queries where the top-ranked candidate is relevant
  - 1525 • **Recall@10 (R@10):** Proportion of queries where at least one relevant item appears in the top-10 results
  - 1526 • **Recall@100 (R@100):** Proportion of queries where at least one relevant item appears in the top-100 results

#### 1529 E.11.3 RETRIEVAL TASK ILLUSTRATION

1531 Figure 6 illustrates the zero-shot retrieval evaluation process. Given a query protein  $p_j$ , the model  
1532 computes cosine similarities with all candidate drugs in the dataset and ranks them by similarity  
1533 scores. Retrieval metrics (R@1, R@10, R@100) measure whether known positive drug-target in-  
1534 teractions appear within the top-k ranked candidates.



1546 Figure 6: Illustration of zero-shot retrieval evaluation. A query protein  $p_j$  is compared against  
1547 all candidate drugs  $\{d_{i-1}, d_i, d_{i+1}, \dots\}$  using cosine similarity of learned embeddings. Recall@k  
1548 metrics evaluate whether any known positive interactions appear in the top-k retrieved candidates.  
1549

#### 1550 E.11.4 IMPLEMENTATION DETAILS

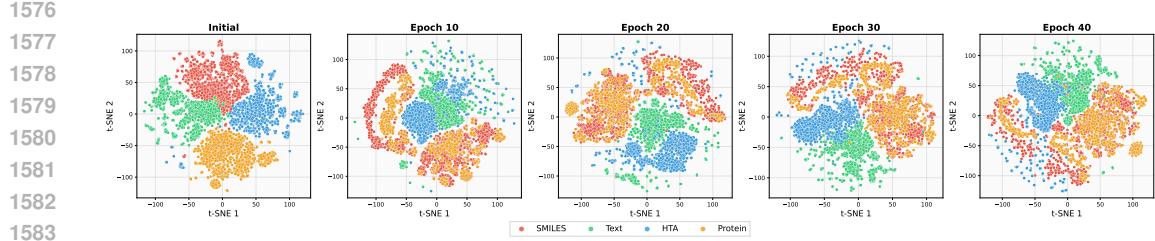
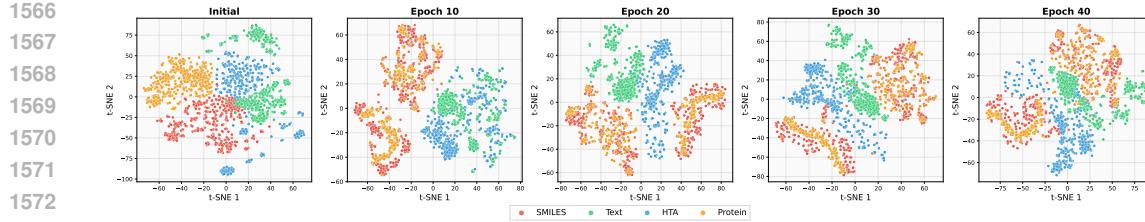
1552 **Model Architecture:** We utilize the same encoder architectures and projector networks as in the  
1553 main pre-training framework:

- 1554 • SMILES projector:  $768 \rightarrow 768 \rightarrow 512 \rightarrow 512$  (with GELU, LayerNorm, Dropout)
- 1555 • Protein projector:  $1280 \rightarrow 768 \rightarrow 512 \rightarrow 512$  (with GELU, LayerNorm, Dropout)

1557 **Batch Processing:** Due to computational constraints, embeddings are generated in batches of 16  
1558 sequences to manage memory usage while maintaining efficiency.  
1559

1560 **No Additional Training:** Critically, no additional training or fine-tuning is performed for the re-  
1561 trieval task. We use the representations learned during the multimodal pre-training phase directly,  
1562 demonstrating the quality of the learned representations.

1563 **Evaluation Protocol:** Following standard practice in information retrieval, we compute metrics  
1564 across all queries in each dataset and report average performance. The evaluation uses only positive  
1565 interactions from the retrieval datasets, ensuring fair assessment of the model's ability to identify  
true drug-target relationships.



1586  
1587  
1588 The strong performance of GRAM-DTI in this zero-shot setting (Table 2 in main text) demonstrates  
1589 that our volume-based multimodal alignment successfully learns semantically meaningful repres-  
1590 entations that capture drug-target relationships without task-specific supervision.

## F COMPREHENSIVE MULTIMODAL EMBEDDING EVOLUTION ANALYSIS

1595 This section provides a comprehensive analysis of how GRAM-DTI learns unified multimodal rep-  
1596 resentations across different sample sizes and training epochs. We examine embedding evolution  
1597 patterns to understand the dynamics of volume-based multimodal alignment and validate the effec-  
1598 tiveness of our adaptive modality dropout mechanism.

### F.1 EXPERIMENTAL SETUP

1600 We conducted embedding evolution analysis across multiple scales to ensure robustness of our ob-  
1601 servations:

- 1602 • **Sample sizes:** 500, 3,000, and 5,000 randomly selected quadruplets
- 1603 • **Training epochs:** Initial state (epoch 0), 10, 20, 30, and 40
- 1604 • **Visualization method:** t-SNE with perplexity=30, max\_iter=1000
- 1605 • **Preprocessing:** L2 normalization of projected embeddings, standardization per modality

1606 For each epoch, we extracted embeddings from the four modalities using their respective pre-  
1607 trained encoders (MolFormer-XL for SMILES, Molt5 for Text/HTA, ESM-2 for Protein), applied  
1608 the trained projection layers to map into the unified 512-dimensional space, and performed t-SNE  
1609 visualization.

## G LARGE LANGUAGE MODELS USAGE STATEMENT

1610 We only used Large Language Models to correct grammars and polish the writing.

