

1	<b>Contents</b>	
2	<b>A Framework</b>	<b>1</b>
3	A.1 Details of Datasets and Instructions . . . . .	1
4	A.2 Details of Framework Application . . . . .	5
5	A.3 Baselines Evaluation . . . . .	5
6	<b>B Method</b>	<b>6</b>
7	B.1 Discussion of Individual Encoding Module Method . . . . .	6
8	B.2 Model Theoretical Capacity . . . . .	7
9	B.3 Detailed Related Work . . . . .	8
10	<b>C Experiments</b>	<b>9</b>
11	C.1 Experiment setting . . . . .	9
12	C.2 Detailed Zero-Shot Result . . . . .	9
13	C.3 Detailed Few-Shot Results . . . . .	11
14	C.4 Detailed Ablation Results of Pretraining . . . . .	13
15	C.5 Instruction Robustness . . . . .	13
16	C.6 Instruction Ablation . . . . .	15
17	C.7 Attention Visualization . . . . .	15
18	<b>A Framework</b>	
19	<b>A.1 Details of Datasets and Instructions</b>	

Table 1: Data Overview

Splitting	Data Class	Dataset	No. of Molecules	No. of Tasks	Task Metric	Task Type
Pretraining	Bioactivity assay	ChEMBL bioassay activity dataset	365065	1048	ROC_AUC	Classification
	Physico-chemical	CHEMBL Property	365065	13	RMSE	Regression
Large Scale	PCBA PubChem HTS bioAssay		437929	128	ROC_AUC	Classification
	ChEMBL Zero-Shot		91266	262	ROC_AUC	Classification
Pharmacokinetic	CYP inhibition		16896	5	ROC_AUC	Classification
	BBBP Blood-brain barrier penetration		2039	1	ROC_AUC	Classification
Downstream Zero-Shot	MUV PubChem bioAssay		93087	17	ROC_AUC	Classification
	BACE-1 benchmark set		1513	1	ROC_AUC	Classification
	HIV replication inhibition		41127	1	ROC_AUC	Classification
	Tox21 Toxicology in the 21st century		7831	12	ROC_AUC	Classification
	Toxcast		8598	617	ROC_AUC	Classification
	ESOL Water solubility		1128	1	RMSE	Regression
Physico-chemical	FreeSolv Solvation free energy		642	1	RMSE	Regression
	Lipo Lipophilicity		4200	1	RMSE	Regression

The datasets used in our study are presented in Table 1. These datasets consist of different types of tasks related to molecule property prediction. It should be noted that during the pretraining phase, the loss function is not specific to the task types, but rather encompasses the generative loss of the language model.

We have chosen not to include certain datasets, namely *SIDER* and *ClinTox*, in our collection of datasets. The decision was based on the fact that the tasks associated with these datasets are not clearly defined and involve complex systemic phenomena, making it challenging to describe them through instructional texts. For instance, the *ClinTox* dataset involves determining whether drugs have passed the FDA approval, which is not an objective problem but rather a dynamic and intricate social phenomenon. The *SIDER* dataset focuses on describing the side effects of drugs on system organ classes, which have intricate mechanisms and a wide range of possible causes, making them difficult to be effectively conveyed through instructions.

32 For the ChEMBL property dataset that we have constructed, detailed information can be found in Table  
33 2. These properties are sourced from the ChEMBL database [19] through the web API.

Table 2: ChEMBL property tasks and labels

Property	Label type
Aromatic rings number	Integer
$\text{cx\_logd}$ distribution coefficient	Real
$\text{cx\_logp}$ partition coefficient	Real
$\text{cx\_most\_apka} - \log_{10}$ dissociation constant	Real
Molecular masses	Real
Hydrogen bond donor number	Integer
Heavy atom number	Integer
Lipinski’s rule of five violation number	Integer
Polar surface area (PSA)	Real
Quantitative Estimate of Druglikeness (QED)	Real
Rule of three passes	Bool
Rotatable bond number	Integer

34 The task explanation is primarily sourced from relevant papers, websites, or databases that introduce  
35 and compile the respective datasets. The specific sources utilized depend on the particular datasets  
36 under consideration. For ChEMBL tasks, we obtain task descriptions from the ChEMBL website.  
37 Descriptions for MoleculeNet tasks and PCBA are primarily sourced from the PubChem website.  
38 Certain datasets, such as Toxcast, include task descriptions within the dataset files. In the case of  
39 other tasks, like ChEMBL property and Physical-Chemical tasks, instructions are derived from Wiki or  
40 other papers. We list the instruction source in Table 3.

Table 3: Data sources and classes for different stages of the model

Dataset	Instruction Source
ChEMBL Zero-Shot bioassay activity dataset	ChEMBL Database
CHEMBL Property	Wiki
PCBA PubChem HTS bioAssay	Pubchem Database
ChEMBL Zero-Shot bioassay activity dataset	ChEMBL Database
CYP PubChem BioAssay CYP 1A2, 2C9, 2C19, 2D6, 3A4 inhibition	Pubchem Database
BBBP Blood-brain barrier penetration	Paper [21]
MUV PubChem bioAssay	Pubchem Database
BACE-1 benchmark set	Pubchem Database
HIV replication inhibition	Paper [38]
Tox21 Toxicology in the 21st century	Pubchem Database
Toxcast	Toxcast file
ESOL Water solubility	Paper [61]
FreeSolv Solvation free energy	Paper [7]
Lipo Lipophilicity	Wiki

41 The description covers a wide range of aspects, including the family, function, and mechanism of  
42 the assay target, the assay experiment setting, the approximation method used for determining the  
43 property, and others. We describe regression tasks by introducing the relationship between the  
44 task property and other properties, i.e. how to estimate these properties by other ones. However,  
45 this method is still challenging due to the model’s capacity to understand complex mathematical  
46 relationships.

47 The instructions for each task are generated automatically by conducting searches on the databases and  
48 summarizing the descriptions. We use a mixture strategy of summarizing, combining template-based  
49 summarizing and GPT-3.5-turbo-based summarizing methods. The GPT-3.5-turbo-based summarizing  
50 method is applied by the prompt ‘Summarize the assay: \n {Descriptions to be summarized}’.

51 The resulting instructions are then concatenated with relevant questions. These instructions are  
52 subsequently reviewed and validated by a professional biology Ph.D. student and slightly modified if  
53 necessary.

54 We then list the instructions of each dataset. For datasets with more than one task, we only list the  
55 instruction of one task as an illustration.

56 Chembl

57 "The assay is PUBCHEM\_BIOASSAY: qHTS Assay for Activators of  
58 Human Muscle isoform 2 Pyruvate Kinase. (Class of assay:  
59 confirmatory) , and it is Direct single protein target  
60 assigned . The assay has properties: assay category is  
61 confirmatory ; assay organism is Homo sapiens ; assay type  
62 description is Functional . Is the molecule effective to this  
63 assay?"

64 Chembl property

65 The partition coefficient, abbreviated P, is defined as a  
66 particular ratio of the concentrations of a solute between the  
67 two solvents (a biphasic of liquid phases), specifically for  
68 un-ionized solutes, and the logarithm of the ratio is thus Log  
69 P. When one of the solvents is water and the other is a  
70 non-polar solvent, then the log P value is a measure of  
71 lipophilicity or hydrophobicity. The defined precedent is for  
72 the lipophilic and hydrophilic phase types to always be in the  
73 numerator and denominator respectively. What is the logarithm  
74 of the partition coefficient of this molecule?

75 PCBA

76 "The assay tests the inhibition of ALDH1A1 activity using  
77 propionaldehyde as an electron donor and NAD+ as an electron  
78 acceptor. The conversion of NAD+ to NADH is measured via an  
79 increase in fluorescence intensity to determine enzyme  
80 activity. ALDH1A1 plays critical roles in the metabolic  
81 activation of retinoic acid and may be a target for inhibitor  
82 development in metabolic diseases. Is the molecule effective  
83 to this assay?"

84 CYP450

85 "Find molecules that can effectively inhibit Cytochrome P450  
86 (CYP450) enzymes, particularly CYP1A2, to help reduce the risk  
87 of adverse drug events and drug-drug interactions caused by  
88 CYP450-mediated metabolic pathways. Consider the various  
89 CYP450 inhibition mechanisms such as occupying active sites or  
90 weakening enzyme activity, while keeping in mind the potential  
91 for increased side effects due to elevated blood drug  
92 concentrations. Is this molecule effective to this assay?"

93 BBBP

94 "In general, molecules that passively diffuse across the brain  
95 blood barrier have the molecular weight less than 500, with a  
96 LogP of 2-4, and no more than five hydrogen bond donors or  
97 acceptors. Does the molecule adhere to the three rules or not?"

98 MUV

99 "Protein kinase A (PKA) is an ubiquitous serine/threonine  
100 protein kinase and belongs to the AGC kinase family. It has  
101 several functions in the cell, including regulation of immune  
102 response, transcription, cell cycle and apoptosis. PKA is a  
103 cAMP dependent enzyme that exists in its native inactive form

104 as a 4 subunit enzyme with two regulatory and two catalytic  
105 subunits. Binding of cAMP to the regulatory subunit leads to  
106 the disassembly of the complex and release of now active  
107 catalytic subunits. Is this molecule inhibitor of PKA?"

#### 108 BACE

109 "BACE1 is an aspartic-acid protease important in the  
110 pathogenesis of Alzheimer's disease, and in the formation of  
111 myelin sheaths. BACE1 is a member of family of aspartic  
112 proteases. Same as other aspartic proteases, BACE1 is a  
113 bilobal enzyme, each lobe contributing a catalytic Asp  
114 residue, with an extended active site cleft localized between  
115 the two lobes of the molecule. The assay tests whether the  
116 molecule can bind to the BACE1 protein. Is this molecule  
117 effective to the assay?"

#### 118 HIV

119 "Human immunodeficiency viruses (HIV) are a type of  
120 retrovirus, which induces acquired immune deficiency syndrome  
121 (AIDs). Now there are six main classes of antiretroviral  
122 drugs for treating AIDs patients approved by FDA, which are  
123 the nucleoside reverse transcriptase inhibitors (NRTIs), the  
124 non-nucleoside reverse transcriptase inhibitors (NNRTIs), the  
125 protease inhibitors, the integrase inhibitor, the fusion  
126 inhibitor, and the chemokine receptor CCR5 antagonist. Is  
127 this molecule effective to this assay?"

#### 128 Tox21

129 "Estrogen receptor alpha (ER aplha) is Nuclear hormone  
130 receptor. The steroid hormones and their receptors are  
131 involved in the regulation of eukaryotic gene expression and  
132 affect cellular proliferation and differentiation in target  
133 tissues. Ligand-dependent nuclear transactivation involves  
134 either direct homodimer binding to a palindromic estrogen  
135 response element (ERE) sequence or association with other  
136 DNA-binding transcription factors, such as AP-1/c-Jun, c-Fos,  
137 ATF-2, Sp1 and Sp3, to mediate ERE-independent signaling. Is  
138 this molecule effective to this assay?"

#### 139 Toxcast

140 "APR\_HepG2\_CellCycleArrest\_24hr, is one of 10 assay  
141 component(s) measured or calculated from the APR\_HepG2\_24hr  
142 assay. It is designed to make measurements of cell phenotype,  
143 a form of morphology reporter, as detected with fluorescence  
144 intensity signals by HCS Fluorescent Imaging technology. Data  
145 from the assay component APR\_HepG2\_CellCycleArrest\_24hr was  
146 analyzed into 2 assay endpoints. \nThis assay endpoint,  
147 APR\_HepG2\_CellCycleArrest\_24h\_dn, was analyzed in the negative  
148 fitting direction relative to DMSO as the negative control and  
149 baseline of activity. \nUsing a type of morphology reporter,  
150 measures of all nuclear dna for loss-of-signal activity can be  
151 used to understand the signaling at the pathway-level as they  
152 relate to the gene . \nFurthermore, this assay endpoint can be  
153 referred to as a primary readout, because this assay has  
154 produced multiple assay endpoints where this one serves a  
155 signaling function. \nTo generalize the intended target to  
156 other relatable targets, this assay endpoint is annotated to

157 the \"cell cycle\" intended target family, where the subfamily  
158 is \"proliferation\". Is this molecule effective to this  
159 assay?"

160 ESOL

161 "Solubility (logS) can be approximated by negative LogP -0.01  
162 \* (Mpt  $\times 25$ ) + 0.5 . Can you approximate the logS of this  
163 molecule by its negative logP and Mpt?"

164 FreeSolv

165 "The free energy of hydration can be approximated by  
166  $\Delta G_{\text{hyd}} = \Delta G_{\text{solv, soln}} - \Delta G_{\text{solv, gas}} + RT \ln$   
167  $(10^{-\text{pKa}})$ . Can you tell me the free energy of hydration (by  
168 using the negative pka) of this molecule, predicted by using  
169  $\Delta G_{\text{solv}}$  and negative pka?"

170 Lipo

171 "Lipophilicity is an important feature of drug molecules that  
172 affects both membrane permeability and solubility, measured by  
173 octanol/water distribution coefficient (logD at pH 7.4).  
174 What's the octanol/water distribution coefficient (logD at pH  
175 7.4) of this molecule?"

## 176 A.2 Details of Framework Application

177 In our framework, we represent the labels of various tasks as strings. For assay tasks involving  
178 classification, the labels are converted to either "Yes" or "No" based on whether the molecule has  
179 an effect on the assay. In regression tasks, the labels are transformed into numerical strings. Integer  
180 values remain unchanged, while decimal numbers are rounded to two decimal places.

181 To conduct zero-shot testing on our model, we generate output sequences and extract the answer from  
182 the results. For assay classification, we consider the first token generated as the answer and use the  
183 scores for the 'Yes' and 'No' tokens to compute the ROC-AUC score for classification. In regression  
184 tasks, we extract the number from the generated sequence by performing string matching using a  
185 regular expression template:  $r"-\d+\.\d*e?\d*?"$ . Notably, we discovered that GIMLET consistently  
186 generates results in the correct format for all classification tasks and accurately formatted numbers  
187 for over 98% of regression testing samples, without any augmentation of restriction in the vocabulary.

## 188 A.3 Baselines Evaluation

189 For the baselines, we apply our instruction-based molecule zero-shot learning to their respective  
190 settings. KVPLM employs SMILES for molecule representation and utilizes masked language  
191 modeling for molecule-text data. Galactica also represents molecules using SMILES but generates  
192 the next sentence in an autoregressive manner. MoMu employs contrastive learning between the  
193 GNN-encoded molecule and the corresponding text, allowing it to score each candidate sentence for  
194 the target molecule and retrieve the best matching one. Our application of each baseline model aligns  
195 with their intended use.

196 It is important to note that for the baseline models, to avoid baselines generating answers in classi-  
197 fication not in our parsing method ('Yes' and 'No'), we limit the vocabulary during generation to  
198 only include 'Yes' and 'No' in classification tasks. This restriction is achieved by utilizing the bias  
199 term in huggingface to prevent the generation of other words. However, it is worth mentioning that  
200 our model, GIMLET, does *not* require this augmentation and is able to generate the desired outputs  
201 *without* any additional constraints.

202 For KVPLM, we mask the answer position in the whole sentence for the model to predict. For  
203 example, for molecule CCOc1cccc1-n1nnc1SCC(=O)NC(=O)NCc1cccc1 and classification tasks  
204 ARE inhibitor, input to KVPLM is:

205 "CCOc1cccc1-n1nnnc1SCC(=O)NC(=O)NCc1ccco1  
206 Oxidative stress has been implicated in the pathogenesis of a  
207 variety of diseases ranging from cancer to neurodegeneration.  
208 The antioxidant response element (ARE) signaling pathway is  
209 important in the amelioration of oxidative stress. Is this  
210 molecule agonists of antioxidant response element (ARE)  
211 signaling pathway? [MASK]"

212 For Galactica, the answer is expected to be generated after reading the question. The input example is

213 "[START\_I\_SMILES] CCOc1cccc1-n1nnnc1SCC(=O)NC(=O)NCc1ccco1  
214 [END\_I\_SMILES]  
215 Question: Oxidative stress has been implicated in the  
216 pathogenesis of a variety of diseases ranging from cancer to  
217 neurodegeneration. The antioxidant response element (ARE)  
218 signaling pathway is important in the amelioration of  
219 oxidative stress. Is this molecule agonists of antioxidant  
220 response element (ARE) signaling pathway?  
221 Answer:"

222 For MoMu, we compute the matching score between the molecule graph and the instruction with  
223 each answer. In the example, the classification scores for 'Yes' and 'No' are computed by matching  
224 graph feature of molecule CCOc1cccc1-n1nnnc1SCC(=O)NC(=O)NCc1ccco1 with

225 "Oxidative stress has been implicated in the pathogenesis of a  
226 variety of diseases ranging from cancer to neurodegeneration.  
227 The antioxidant response element (ARE) signaling pathway is  
228 important in the amelioration of oxidative stress. Is this  
229 molecule agonists of antioxidant response element (ARE)  
230 signaling pathway? Yes"

231 and

232 "Oxidative stress has been implicated in the pathogenesis of a  
233 variety of diseases ranging from cancer to neurodegeneration.  
234 The antioxidant response element (ARE) signaling pathway is  
235 important in the amelioration of oxidative stress. Is this  
236 molecule agonists of antioxidant response element (ARE)  
237 signaling pathway? No"

238 .

## 239 B Method

### 240 B.1 Discussion of Individual Encoding Module Method

241 The Individual encoding module-based multimodal language model can be formalized as  
242  $LLM(M(G), T)$ , where  $M$  is the individual encoding module for graph data  $G$ . For example,  
243 the visual module is applied to pre-encode the image data to get the dense representation, then put  
244 into the language model as tokens embedding [4, 9, 1, 28]. Current works on molecule language  
245 models also use a GNN to get the representation of molecules to interact with the language models  
246 [16, 49, 48].

247 This method can be considered as decomposition of the conditional probability  $P(y|G, T)$

$$P(\hat{y}|G, T) = \int P_M(z|G)P_{LLM}(\hat{y}|z, T)dz, \quad (1)$$

248 based on the assumption that the feature distributions  $P(z|G)$  should be modeled by modality-specific  
249 modules to introduce inductive bias, and be independent of text information to help with adaptation  
250 to novel text data.

251 However, for the molecule-text model, individual pre-encoding modules present problems. First,  
 252 graph learning relies on structure information, but the dense vectors encoded by GNN have a limited  
 253 capacity to carry structure information, and language models don't have inductive bias toward graph  
 254 structure. Furthermore, training the additional module is difficult due to the increased layers, since  
 255 deep transformers have vanishing gradients in early layers [29, 2], which is a well-known problem of  
 256 transformer. Lastly, the additional modules increase parameters and training costs.

257 Our method GIMLET not only overcome these issues, our approach GIMLET not only directly unifies  
 258 the standard language model for graph and text *without* introducing additional graph encoder module,  
 259 but also remains the decoupled graph encoding for better generalization.

## 260 B.2 Model Theoretical Capacity

261 In this section, we analyze the theoretical capacity of our modeling method.

262 **Theorem 1** *Assume for different input features and position embeddings, the transformer layers can*  
 263 *output different output features. The transformer with distance-based relative position embedding*  
 264 *has a stronger capacity than the 1-WL test for the graph isomorphism problem.*

265 **Proof 1** *The 1-WL test is defined as the following iteration:*

$$\begin{aligned} \chi_G^0(i) &= \text{hash}(v_i) \\ \chi_G^t(i) &:= \text{hash}(\chi_G^{t-1}(i), \{\chi_G^{t-1}(j) : j \in \mathcal{N}_G(i)\}) \quad (\forall i \in N), \end{aligned} \tag{2}$$

266 where  $\chi_G$  is the label in WL test, hash is the hash function,  $\mathcal{N}_G(i)$  is the neighbor of node  $i$ .

267 The transformer with distance-based relative position embedding can be considered as the following  
 268 mapping:

$$\begin{aligned} \chi_G^t(i) &:= \text{hash}(\{(d_G(i, j), \chi_G^{t-1}(j)) : j \in N\}) \\ &= \text{hash}(\{(0, \chi_G^{t-1}(i))\} \\ &\cup \{(1, \chi_G^{t-1}(j)) : j \in \mathcal{N}_G(i)\} \\ &\cup \{(d_G(i, k), \chi_G^{t-1}(k)) : k \in N - \mathcal{N}_G(i) - \{i\}\}) \end{aligned} \tag{3}$$

269 It can be seen that the iteration of the transformer with distance-based relative position embedding  
 270 includes both the node  $i$  and its neighbors  $\mathcal{N}_G(i)$ , marked by distance 0 and 1, respectively, ensuring  
 271 the capacity is at least as strong as 1-WL test. It further includes other nodes far away, along with  
 272 their distance, which constitutes a stronger capacity than 1-WL test. Figure 1 are two example  
 273 graphs that cannot be distinguished by 1-WL test, but can be distinguished by transformer with  
 274 distance-based relative position embedding.

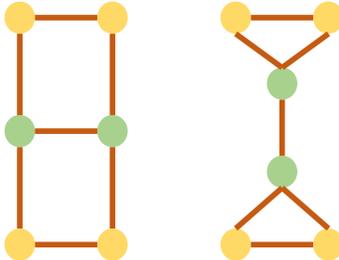


Figure 1: Two example graphs that cannot be distinguished by 1-WL test, but can be distinguished by transformer with distance-based relative position embedding.

275 **Theorem 2** *Assume for different input features and position embeddings, the transformer layers can*  
 276 *output different output features. GIMLET can distinguish graph-instruction pairs if graphs can be*

277 *distinguished by transformer with distance-based relative position embedding, or instructions are*  
278 *different.*

279 **Proof 2** *As GIMLET decomposes the attention from graph nodes to text, the graph nodes can only*  
280 *attend to other graph nodes. Thus the encoding capacity of graph data is the same as a single*  
281 *transformer with distance-based relative position embedding for graph data.*

282 *Along with the assumption of transformer layers, GIMLET is able to distinguish graph-instruction*  
283 *pairs if graphs can be distinguished by transformer with distance-based relative position embedding,*  
284 *or instructions are different.*

### 285 **B.3 Detailed Related Work**

286 We present a detailed related work here, due to the space limitation of paper.

287 **Molecule Representation learning** In recent years, there has been a growing interest in developing  
288 molecular representation learning for downstream tasks like drug discovery and other applications.  
289 One approach that has received considerable attention is utilizing language modeling techniques to  
290 acquire molecular representations based on Simplified Molecular Input Line Entry System (SMILES)  
291 strings [57, 10]. Although sequence-based representations have demonstrated success in some ap-  
292 plications, concerns have been raised about their capability to incorporate all pertinent substructure  
293 information. To address this limitation, some researchers have proposed the use of Graph Neural Net-  
294 works (GNNs) to model molecules as graphs [20, 67, 25], potentially providing a more comprehensive  
295 and accurate representation of the molecular structure.

296 Existing GNNs follow the message-passing paradigm and suffer from problems like long-range  
297 dependency vanishing and over-smoothing. Recently, Graph Transformer [44, 65] has been proposed  
298 to better encode structures of graphs. The Graph Transformer is inspired by the Transformer  
299 architecture, which has shown remarkable performance in natural language processing [55, 13, 33].  
300 The Graph Transformer extends the Transformer architecture to the graph domain, allowing the  
301 model to capture the global structure and long-range dependencies of the graph [69, 14, 27, 26, 62,  
302 40, 34, 65, 8, 35, 11, 5, 22, 71].

303 **Molecule Pretraining** To fully explore the inherent structural information of molecules on a large  
304 scale and transfer useful information to downstream tasks, significant efforts have been made to  
305 address the inadequacies in molecular pre-training. Supervised pretraining is commonly used for  
306 learning useful representations [25, 65, 52]. As for unsupervised pretraining, one approach involved  
307 using an generative pre-training strategy on molecular SMILES strings [57, 24, 10, 3, 45] and Graph  
308 [25, 30, 44, 70], which was followed by recent works adopting the contrastive paradigm that aligns  
309 representation of augmented views of the same graph but keeping views from other graphs away  
310 [56, 50, 23, 67, 66, 53, 64, 18, 51, 59, 63, 58, 32].

311 The pretraining methods mentioned focus on obtaining representations for supervised training.  
312 However, for natural language instruction-based zero-shot graph learning, it’s necessary to incorporate  
313 natural language into the pretraining process. Several studies have explored molecule structure-text  
314 multimodal pretraining. One class of method is the SMILES based language model, including  
315 KVPLM [68] and MolT5 [15], which use SMILES strings and text for joint representation and  
316 translation. Another work Galactica [54] explored the multi-task molecule task learning with  
317 instruction. Some other works acquire advanced representations for molecules by GNN, such as  
318 Text2Mol [16], MoMu [49], MoleculeSTM [31], and CLAMP [48], trained by contrastive learning  
319 between molecule graph and text description for molecule retrieval and caption tasks. MoleculeSTM  
320 and CLAMP explored molecule editing and property prediction with instructions. However, none of  
321 these works address the zero-shot fashion on complex molecule tasks like property prediction, due to  
322 constraints imposed by the pretraining methodology that not addressing the instruction-following  
323 ability, and their model capacity for representing molecule graphs.

324 **Instruction-based zero-shot learning** Instruction-based zero-shot learning is an innovative approach  
325 that leverages natural language instructions and definitions to enable neural models to solve a variety  
326 of tasks [42, 6, 47, 17, 72, 36, 37, 41]. By providing a human-readable prompt, this method enables  
327 easier and more efficient specification of the learning task by utilizing knowledge about the task  
328 without data. To enhance the model’s ability to follow instructions, some researchers have employed  
329 instruction-based pretraining techniques [46, 60, 12, 39], which explicitly train language models

330 to solve tasks with instructions. Besides natural language processing, instruction-based zero-shot  
331 learning is also studied in multimodal domains like images [4, 9, 1, 28].

## 332 C Experiments

### 333 C.1 Experiment setting

334 Our model only utilizes the basic features [25, 52] of molecule graphs, which do not include additional  
335 features like ring markers. Specifically, it utilizes the first two dimensions of node features and the  
336 first two dimensions of edge features processed by ogb.smiles2graph. Therefore, the effectiveness  
337 of GIMLET predominantly stems from its architectural design and pretraining rather than the graph  
338 features it incorporates.

339 Following the standard supervised setting in previous studies [25], we utilize the scaffold strategy  
340 [43] to partition datasets into three subsets: the training set, validation set, and testing set with a  
341 ratio of 0.8, 0.1, 0.1. The scaffold strategy is a deterministic approach that involves sorting the data  
342 based on the scaffold, which represents the molecular structure. While this strategy aids in dataset  
343 partitioning, it can introduce a significant domain gap between the training and testing sets, thereby  
344 increasing the challenge of generalization.

345 For zero-shot, we report the results on the testing sets, ensuring the comparability of our results to  
346 previous works. For few-shot, we report the result of the best validation model on the testing set, the  
347 same as previous works and other supervised baselines [43].

348 Many datasets encompass multiple tasks. To evaluate these datasets, we conduct separate testing for  
349 each task, accompanied by their respective instructions. For datasets with multiple tasks, we report  
350 the average ROC-AUC score for each task, following the methodology established in previous works  
351 [25].

### 352 C.2 Detailed Zero-Shot Result

353 We list the full zero-shot result of GIMLET and baselines in Table 4, 5, and 6. The standard deviation  
354 for supervised results are denoted after  $\pm$ , and the multi-task setting results of Galactica are denoted  
355 in parentheses with italic. We also include the instruction-based zero-shot result reported in recent  
356 baseline CLAMP [48] which is tested by their instruction, denoted by italics too. CLAMP is a  
357 contrastive pretrained model with ensembled encoders for molecule and text. The parameter number  
358 for CLAMP’s result is not clearly stated in their paper but should be larger than 10B as they use sT5  
359 language model [42] XXL variant (11B) as one of the ensembled language models.

Table 4: Zero shot performance over Bio-activity tasks

Method	Parameter	Type	bace	hiv	muv	Avg. bio
KVPLM	110M		0.5126	0.6120	0.6172	0.5806
MoMu	113M	Zero Shot	0.6656	0.5026	0.6051	0.5911
CLAMP	> 10B		<i>0.6476</i>	<i>0.8067</i>	-	-
GIMLET	64M		0.6957	0.6624	0.6439	0.6673
Galactica-125M	125M	Multi Task	0.4451( <i>0.561</i> )	0.3671( <i>0.702</i> )	0.4986	0.4369
Galactica-1.3B	1.3B		0.5648( <i>0.576</i> )	0.3385( <i>0.724</i> )	0.5715	0.4916
GCN	0.5M		<i>0.736±0.030</i>	<i>0.757±0.011</i>	<i>0.732±0.014</i>	0.742
GAT	1.0M		<i>0.697±0.064</i>	<i>0.729±0.018</i>	<i>0.666±0.022</i>	0.697
GIN	1.8M	Supervised	<i>0.701±0.054</i>	<i>0.753±0.019</i>	<i>0.718±0.025</i>	0.724
Graphormer	48M		0.7760±0.015	0.7452±0.014	0.7061±0.027	0.7424
Graphormer-p	48M		0.8575±0.006	0.7788±0.012	0.7480±0.020	0.7948

360 The result in parentheses represents the outcome of the multitask setting, also referred to as weakly  
361 supervised in the original paper, where the same instructions are used for both pretraining and testing.  
362 While Galactica has been exposed to the same task instructions, it actually employs multitask learning  
363 with instructions serving as task identity.

364 Even in comparison to Galactica’s multitask result, GIMLET demonstrates comparable or superior  
365 performance on most datasets. This highlights the ability of GIMLET to perform zero-shot tasks with  
366 high quality.

Table 5: Zero shot performance over Toxicity tasks

Method	Parameter	Type	tox21	toxcast	Avg. tox
KVPLM	110M	Zero Shot	0.4917	0.5096	0.5007
MoMu	113M		0.5757	0.5238	0.5498
CLAMP	> 10B		0.6058	0.5383	0.5721
GIMLET	64M		0.6119	0.5904	0.6011
Galactica-125M	125M	Multi Task	0.4964(0.543)	0.5106(0.518)	0.5035
Galactica-1.3B	1.3B		0.4946(0.606)	0.5123(0.589)	0.5035
GCN	0.5M	Supervised	0.749±0.008	0.633±0.009	0.691
GAT	1.0M		0.754±0.005	0.646±0.006	0.700
GIN	1.8M		0.740±0.008	0.634±0.006	0.687
Graphormer	48M		0.7589±0.004	0.6470±0.008	0.7029
Graphormer-p	48M		0.7729±0.006	0.6649±0.006	0.7189

Table 6: Zero shot performance over Pharmacokinetic tasks

Method	Parameter	Type	bbbp	cyp450	Avg. pha
KVPLM	110M	Zero Shot	0.6020	0.5922	0.5971
MoMu	113M		0.4981	0.5798	0.5390
CLAMP	> 10B		0.4788	-	-
GIMLET	64M		0.5939	0.7125	0.6532
Galactica-125M	125M	Multi Task	0.6052(0.393)	0.5369	0.5711
Galactica-1.3B	1.3B		0.5394(0.604)	0.4686	0.5040
GCN	0.5M	Supervised	0.649±0.030	0.8041±0.005	0.7266
GAT	1.0M		0.662±0.026	0.8281±0.004	0.7451
GIN	1.8M		0.658±0.045	0.8205±0.012	0.7392
Graphormer	48M		0.7015±0.013	0.8436±0.003	0.7725
Graphormer-p	48M		0.7163±0.009	0.8877±0.004	0.8020

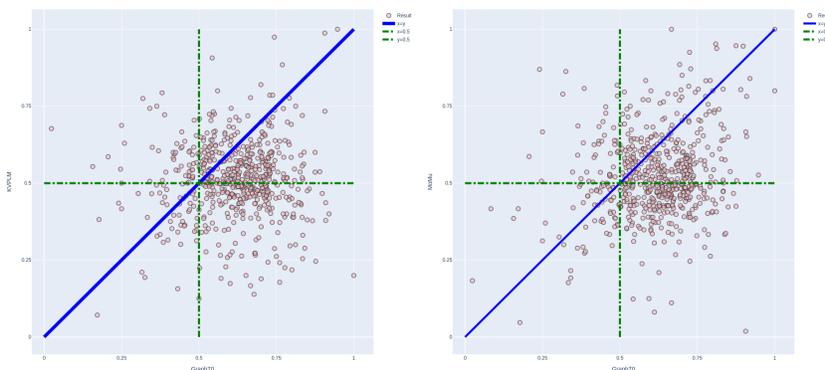


Figure 2: Scatter of GIMLET over baselines. Below the diagonal line  $x=y$  means our method performs better.

367 The disparity between the multitask result and the tested result with our instructions is due to the gap  
368 between their instructions and ours, which indicates that Galactica relies on specific task instructions  
369 for task recognition, without a true understanding of the instructions. As a result, it exhibits poor  
370 generalization to other instruction forms. Note that Galactica even do not surpass KVPLM and  
371 MoMu which are also zero-shot learning methods.

372 GIMLET exhibits superior performance compared to the larger model CLAMP on the majority of  
373 datasets, with the exception of HIV. It is important to highlight that our model is significantly  
374 smaller in size than CLAMP, underscoring the effectiveness of our unified graph-text language model.  
375 Additionally, it should be noted that CLAMP lacks the capability to handle regression tasks due to its  
376 contrastive model architecture, whereas our encoder-decoder architecture enables us to successfully  
377 tackle a wide range of task types.

378 Significantly, the supervised results shed light on the task difficulties associated with each dataset.  
379 This showcases GIMLET’s capability to effectively solve molecule tasks in a zero-shot manner,  
380 approaching the performance of supervised results. Furthermore, our pretraining tasks yield an  
381 average performance improvement of 3 percent for Graphormer, with the largest gains observed in  
382 Bioactivity tasks and the smallest in Toxicity tasks. This suggests that there still exist gaps between  
383 the pretraining data and our downstream tasks, addressing the zero-shot setting of our dataset.

384 In Figure 2, we present scatter plots comparing GIMLET with KVPLM and MoMu across all tasks.  
385 The diagonal line represents the equality line where  $x=y$  indicates our method outperforms the  
386 baseline. Notably, it is evident that GIMLET consistently performs significantly better than random  
387 guessing and surpasses the baselines on all tasks.

388 We plot the scatter of regression tasks in Figure 3. The plot clearly demonstrates a strong correlation  
389 between the predicted and actual values for ESOL and Lipo.

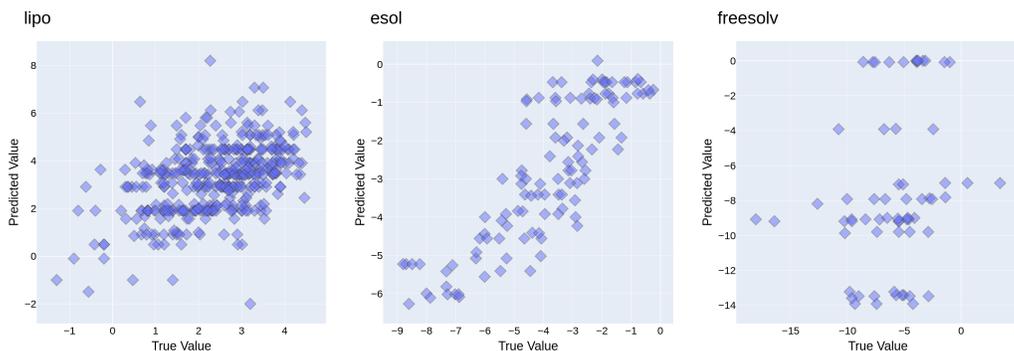


Figure 3: Scatter of GIMLET on generative tasks.

### 390 C.3 Detailed Few-Shot Results

391 In both classification tasks and regression tasks, we fine-tune the last linear layer of all models using  
392 their respective modeling loss.

393 It is important to note that the instruction-based few-shot approach is trained on each task individually,  
394 while supervised baselines are trained on multiple tasks from the dataset. Therefore, comparing  
395 these two approaches may not be strictly fair, as the multitask learning of the supervised baseline can  
396 contribute to improved task performance.

397 The results for few-shot learning on each dataset are presented in Figure 4. It is evident that, across  
398 the majority of datasets, GIMLET demonstrates improvement as the number of few-shot examples  
399 increases. In fact, it even outperforms or matches the performance of the supervised GIN on several  
400 datasets, such as bace, bbbp, and esol. There is also observable enhancement in performance across  
401 various datasets when employing few-shot learning, including tox21, toxcast, lipo, and freesolv.

402 There is not result of MoMu on regression tasks, because MoMu is a contrastive model between  
403 graph and text, which cannot handle regression tasks.

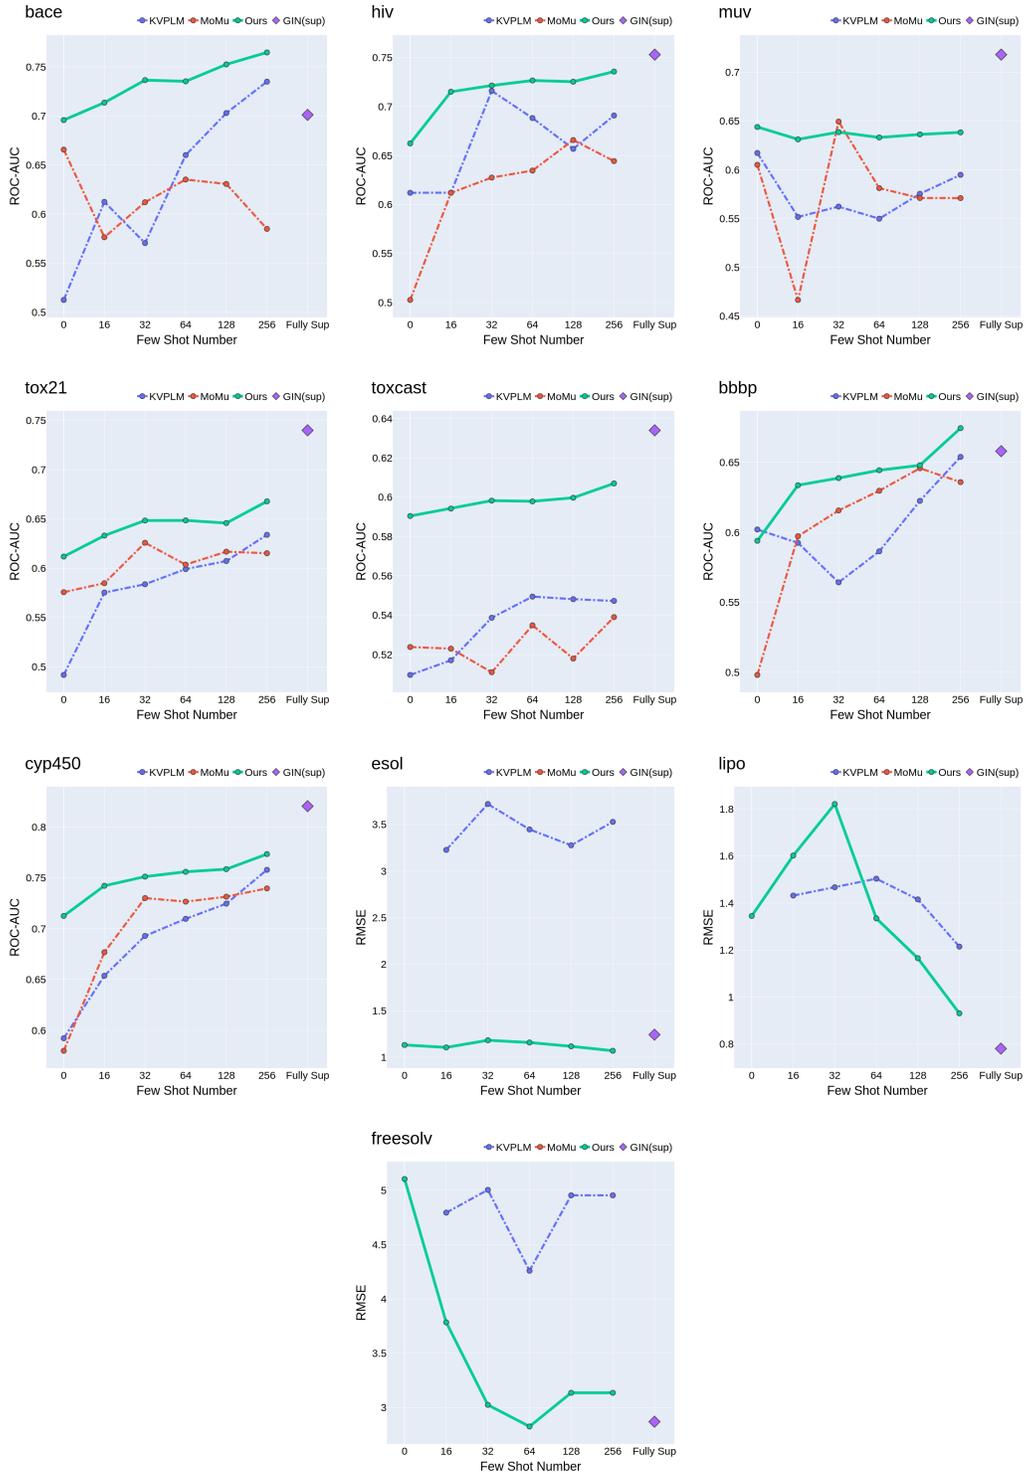


Figure 4: Few-shot performance on each dataset

#### 404 C.4 Detailed Ablation Results of Pretraining

405 The results of pretraining ablation for each dataset are presented in Table 7, 8, 9, and 10. The findings  
406 indicate that both bioactivity assay and physico-chemical properties offer significant benefits for all  
407 the downstream tasks, demonstrating positive transfer across different domains.

Table 7: Pretraining ablation study on Bio-activity tasks

	bace	hiv	muv	Average_bio
bioactivity assay only	0.6390	0.6772	0.6044	0.6402
physico-chemical only	0.4648	0.5461	0.4572	0.4894
both	0.6957	0.6624	0.6439	0.6673

Table 8: Pretraining ablation study on Toxicity tasks

	tox21	toxcast	Average_tox
bioactivity assay only	0.5726	0.5625	0.5676
physico-chemical only	0.4478	0.5017	0.4748
both	0.6119	0.5904	0.6011

Table 9: Pretraining ablation study on Pharmacokinetic tasks

	bbbp	cyp450	Average pha
bioactivity assay only	0.5313	0.6829	0.6071
physico-chemical only	0.5932	0.4976	0.5454
both	0.5939	0.7125	0.6532

#### 408 C.5 Instruction Robustness

409 To test the robustness of GIMLET, the Instructions are rephrased by GPT-3.5-turbo. There are four  
410 types of rephrasing, realized by the following prompts:

411 rewrite

412 'Rephrase the text of the following prompt: \n'

413 expand

414 'Rephrase the text of the following prompt longer: \n'

415 detail

416 'Rephrase the text of the following prompt by adding more  
417 explanation: \n'

418 short

419 'Rephrase the text of the following prompt shorter: \n'

420 Given a task instruction, we rephrase the instruction by the prompts above. Here is an example of  
421 four types of rephrased task instruction from Toxcast:

422 origin

423 "CEETOX\_H295R\_ANDR, is one of 23 assay component(s) measured  
424 or calculated from the CEETOX\_H295R assay. It is designed to  
425 make measurements of hormone induction, a form of inducible  
426 reporter, as detected with absorbance signals by HPLC-MS-MS  
427 technology. Data from the assay component CEETOX\_H295R\_ANDR was  
428 analyzed into 2 assay endpoints. This assay endpoint,  
429 CEETOX\_H295R\_ANDR\_dn, was analyzed in the positive fitting  
430 direction relative to DMSO as the negative control and  
431 baseline of activity. Using a type of inducible reporter,

Table 10: Pretraining ablation study on Physical-chemical tasks

	esol	freesolv	lipo	Average_phy
bioactivity assay only	-	-	-	-
physico-chemical only	1.1822	5.2935	1.3778	2.6178
both	1.1320	5.1027	1.3450	2.5266

432 loss-of-signal activity using HPLC-MS-MS was used to  
 433 understand synthesis of Androstenedione in H295R cell line at  
 434 48hr of chemical exposure. To generalize the intended target  
 435 to other relatable targets, this assay endpoint is annotated  
 436 to the steroid hormone intended target family, where the  
 437 subfamily is androgens. Is this molecule effective to this  
 438 assay?"

439 rewrite

440 "The CEETOX\_H295R\_ANDR is a component of the CEETOX\_H295R  
 441 assay, which measures hormone induction through absorbance  
 442 signals detected by HPLC-MS-MS technology. The data from  
 443 CEETOX\_H295R\_ANDR has been analyzed into two endpoints, with  
 444 CEETOX\_H295R\_ANDR\_dn being analyzed in relation to DMSO as the  
 445 negative control and baseline of activity. HPLC-MS-MS was used  
 446 to detect loss-of-signal activity and understand the synthesis  
 447 of Androstenedione in H295R cell line after 48 hours of  
 448 chemical exposure. This assay endpoint is related to the  
 449 steroid hormone intended target family, specifically the  
 450 subfamily of androgens, and can be generalized to other  
 451 similar targets. Can this assay be effectively performed using  
 452 this molecule?"

453 expand

454 "The CEETOX\_H295R\_ANDR assay component is just one of the 23  
 455 assay components that are measured or calculated from the  
 456 CEETOX\_H295R assay. The assay is specifically designed to  
 457 measure hormone induction, which is a form of inducible  
 458 reporter, and is detected using absorbance signals by means of  
 459 HPLC-MS-MS technology. The data obtained from the  
 460 CEETOX\_H295R\_ANDR assay component was analyzed into two assay  
 461 endpoints. The CEETOX\_H295R\_ANDR\_dn assay endpoint was  
 462 analyzed in the positive fitting direction in relation to DMSO  
 463 as the negative control and activity baseline. To understand  
 464 the synthesis of Androstenedione in the H295R cell line after  
 465 48 hours of chemical exposure, loss-of-signal activity was  
 466 used with HPLC-MS-MS technology. This endpoint is annotated to  
 467 the steroid hormone intended target family to help other  
 468 related targets, where the subfamily is androgens. Can it be  
 469 determined if this particular molecule exhibits desirable  
 470 efficacy to be utilized in this particular assay?"

471 detail

472 "The CEETOX\_H295R\_ANDR is an assay component that is one of  
 473 the 23 components that are measured or calculated from the  
 474 CEETOX\_H295R assay. It is intended to measure hormone  
 475 induction, which is a form of inducible reporter, and the  
 476 measurement is done with the help of absorbance signals using  
 477 HPLC-MS-MS technology. The data obtained from the measurement  
 478 of assay component CEETOX\_H295R\_ANDR is analyzed into two  
 479 assay endpoints. One of these endpoints, CEETOX\_H295R\_ANDR\_dn,

480 is analyzed in the positive fitting direction, relative to  
481 DMSO, which is used as the negative control and baseline for  
482 activity. The HPLC-MS-MS technology is used to detect the  
483 loss-of-signal activity, which helps in understanding the  
484 synthesis of Androstenedione in H295R cell line after 48 hours  
485 of chemical exposure. To make the intended target more  
486 comprehensive and relatable to other targets, the assay  
487 endpoint is annotated to the steroid hormone intended target  
488 family, where the subfamily is androgens. Can this molecule be  
489 used for this assay?"

490 short

491 "CEETOX\_H295R\_ANDR is one of 23 components in the CEETOX\_H295R  
492 assay, measuring hormone induction detected with absorbance  
493 signals by HPLC-MS-MS. It's analyzed into 2 endpoints, with  
494 CEETOX\_H295R\_ANDR\_dn being the positive fitting direction  
495 relative to the negative control. It analyzes the  
496 loss-of-signal activity to understand Androstenedione  
497 synthesis in H295R cell line after 48hr chemical exposure.  
498 It's annotated as a steroid hormone intended target in  
499 androgens sub-family. Is molecule suitable for assay?"

## 500 C.6 Instruction Ablation

501 To ablate the explanation-based instruction, we remove the explanation and only keep the assay name.  
502 The ablated instruction for the instruction above is:

503 "The assay name is CEETOX\_H295R\_ANDR. Is this molecule  
504 effective to this assay?"

## 505 C.7 Attention Visualization

506 We present visualizations of the attention of text tokens to molecule graphs, demonstrating how  
507 our unified transformer incorporates molecule information using various instructions. We randomly  
508 sample molecules and attention heads for visualization. To emphasize high-level features, we focus  
509 on visualizing the attention patterns of the last layer. The redder means the larger attention value.

510 For BACE instruction, we visualize the attention of several keywords marked in red to molecules:

511 "BACE1 is an **aspartic-acid** protease important in the pathogenesis of **Alzheimer's** disease, and in  
512 the formation of **myelin** sheaths. BACE1 is a member of family of **aspartic** proteases. Same as other  
513 aspartic proteases, BACE1 is a **bilobal** enzyme, each lobe contributing a **catalytic** Asp residue, with  
514 an extended active site **cleft** localized between the two lobes of the molecule. The assay tests whether  
515 the molecule can bind to the BACE1 protein. Is this molecule **effective** to the assay?"

516 For BBBP instruction:

517 'In general, molecules that passively diffuse across the **brain blood barrier** have the **molecular weight**  
518 less than **500**, with a **LogP** of **2-4**, and no more than **five hydrogen bond donors** or **acceptors**. Does  
519 the molecule adhere to the **three rules** or not?'

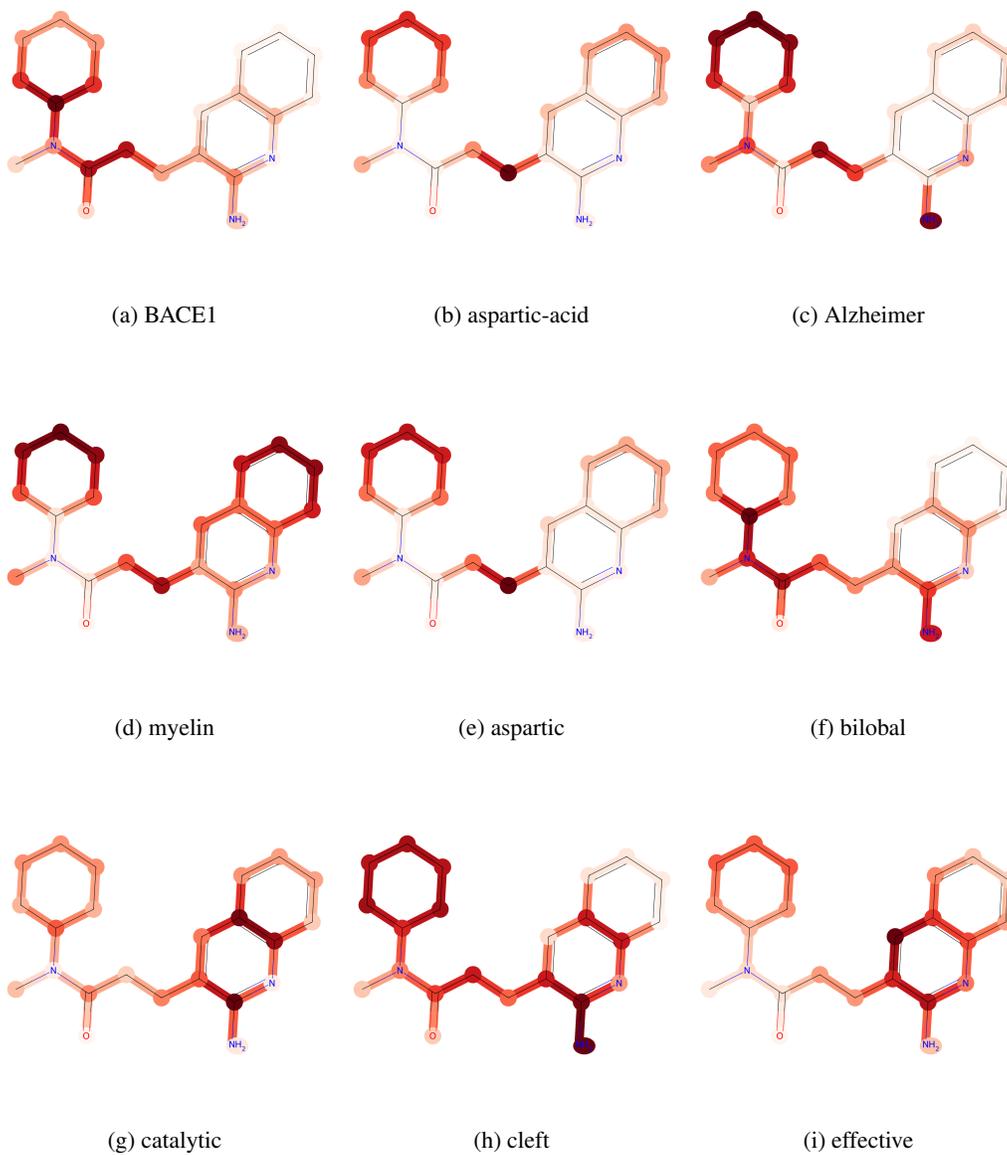


Figure 5: Visualization of attention for BACE on molecule O=C(N(C)C1CCCCC1)CCc1cc2c(nc1N)cccc2



(a) BACE1

(b) aspartic-acid

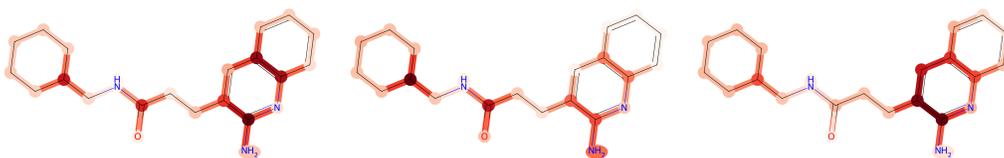
(c) Alzheimer



(d) myelin

(e) aspartic

(f) bilobal



(g) catalytic

(h) cleft

(i) effective

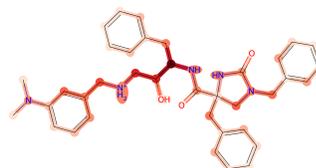
Figure 6: Visualization of attention for BACE on molecule  
O=C(NCC1CCCC1)CCc1cc2c(nc1N)cccc2



(a) BACE1



(b) aspartic-acid



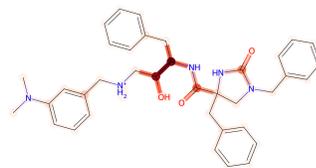
(c) Alzheimer



(d) myelin



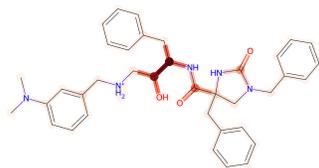
(e) aspartic



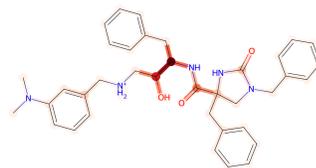
(f) bilobal



(g) catalytic



(h) cleft



(i) effective

Figure 7: Visualization of attention for BACE on molecule

O=C1NC(CN1Cc1cccc1)(Cc1cccc1)C(=O)NC(Cc1cccc1)C(O)C[NH2+][C]c1cc(N(C)C)ccc1

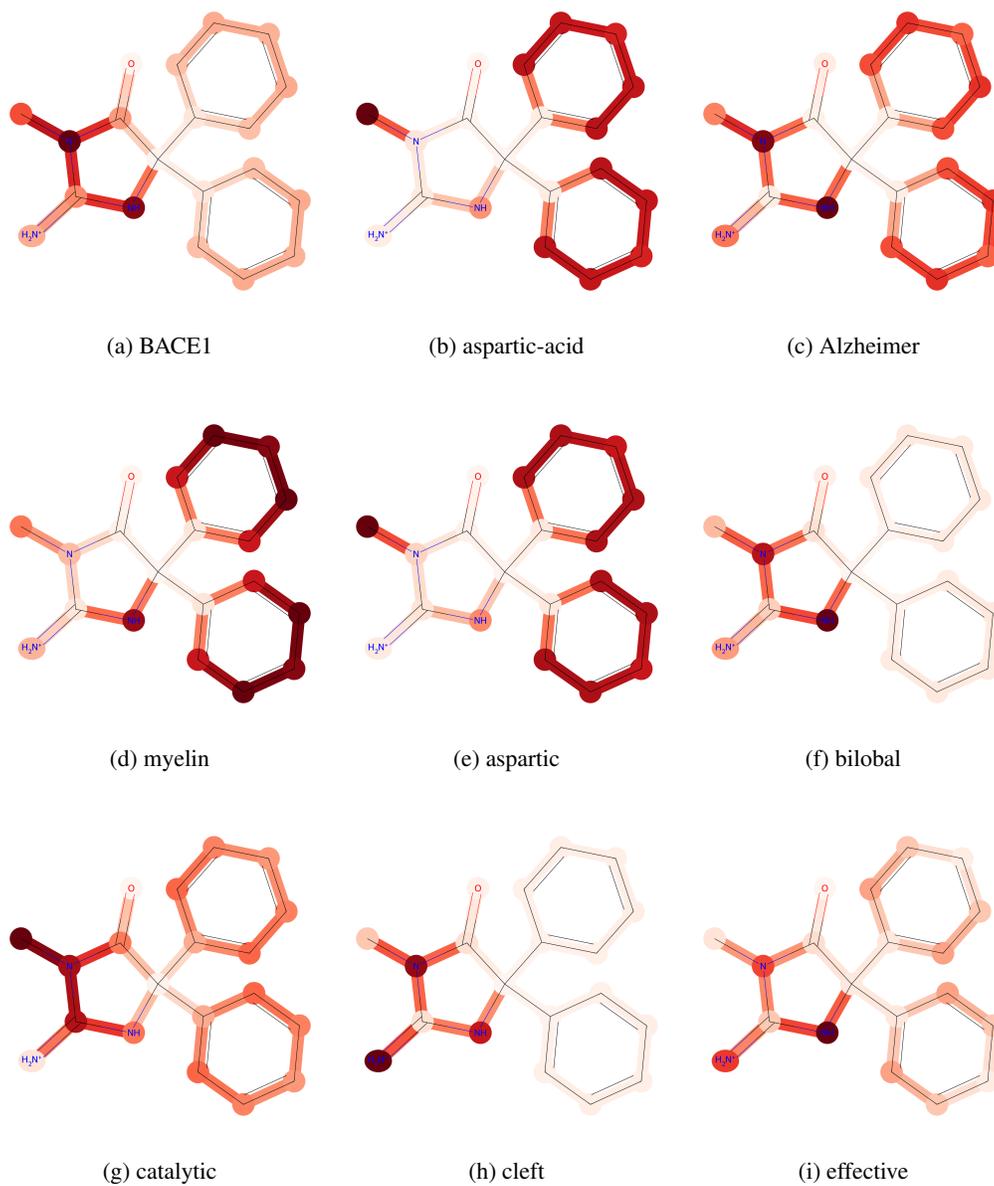
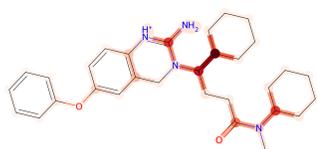


Figure 8: Visualization of attention for BACE on molecule  
O=C1N(C)C(=[NH2+])NC1(c1ccccc1)c1ccccc1



(a) BACE1



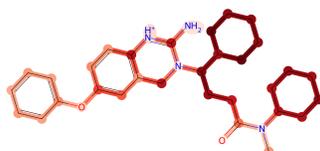
(b) aspartic-acid



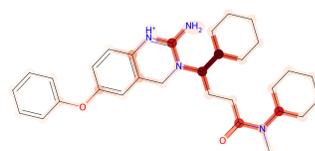
(c) Alzheimer



(d) myelin



(e) aspartic



(f) bilobal



(g) catalytic



(h) cleft



(i) effective

Figure 9: Visualization of attention for BACE on molecule  
O(c1cc2CN(C(CCC(=O)N(C)C3CCCCC3)C3CCCCC3)C(=[NH+]c2cc1)N)c1cccc1

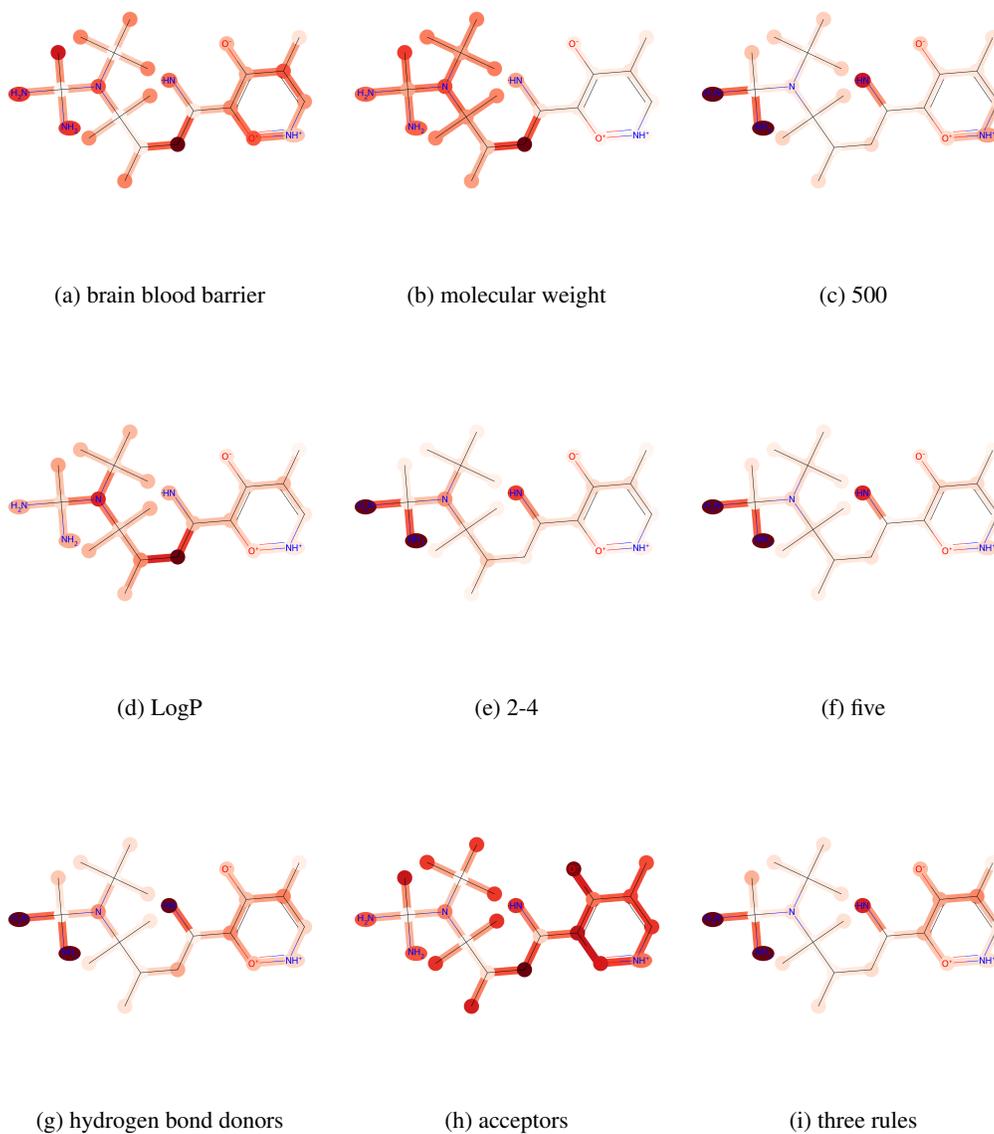


Figure 10: Visualization of attention for BBBP on molecule

C(CC(C)C([N@@](C(C)(C)C)C(N)(C)N)(C)C)c1c(c([nH+][o+])1)C[O-]

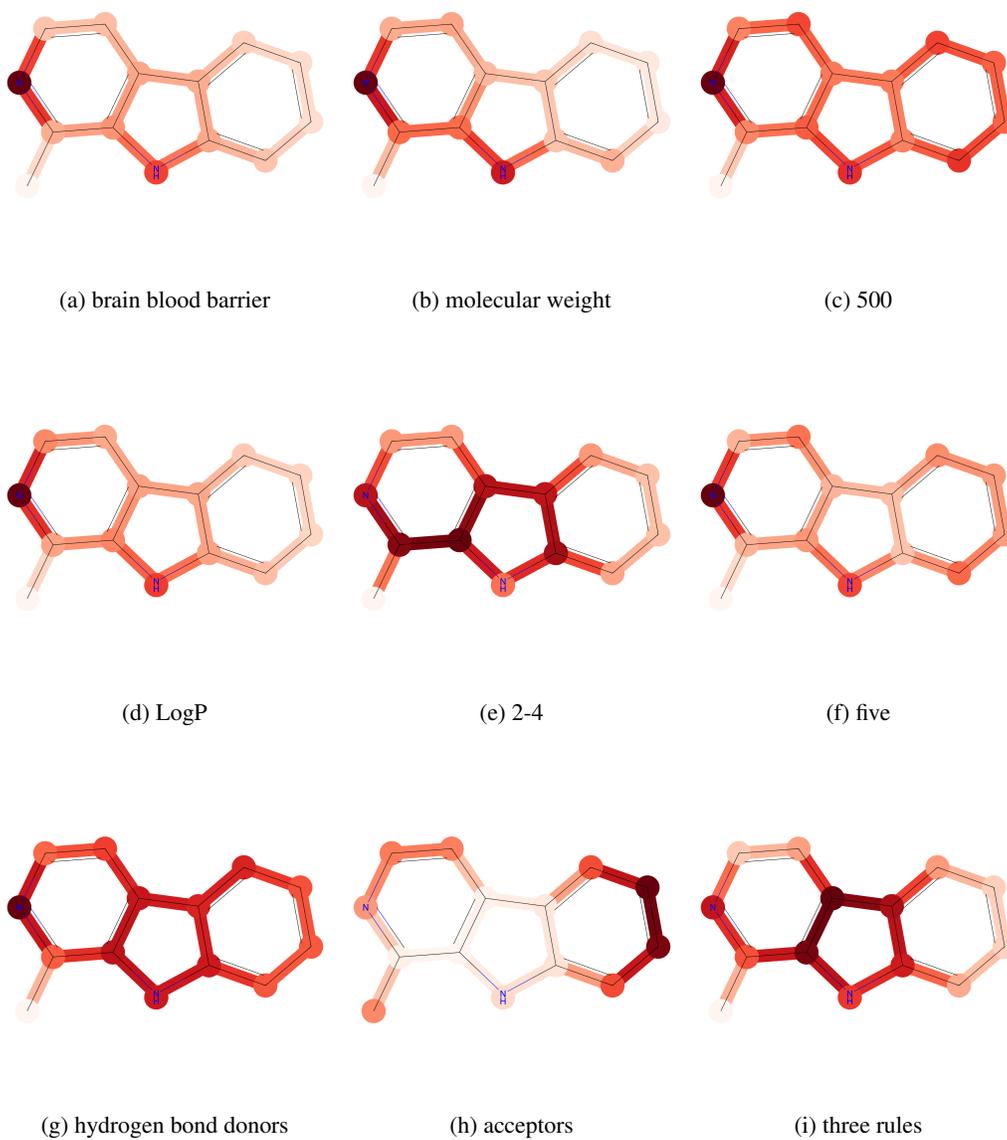


Figure 11: Visualization of attention for BBBP on molecule  
Cc1nccc2c1[nH]c3cccc23

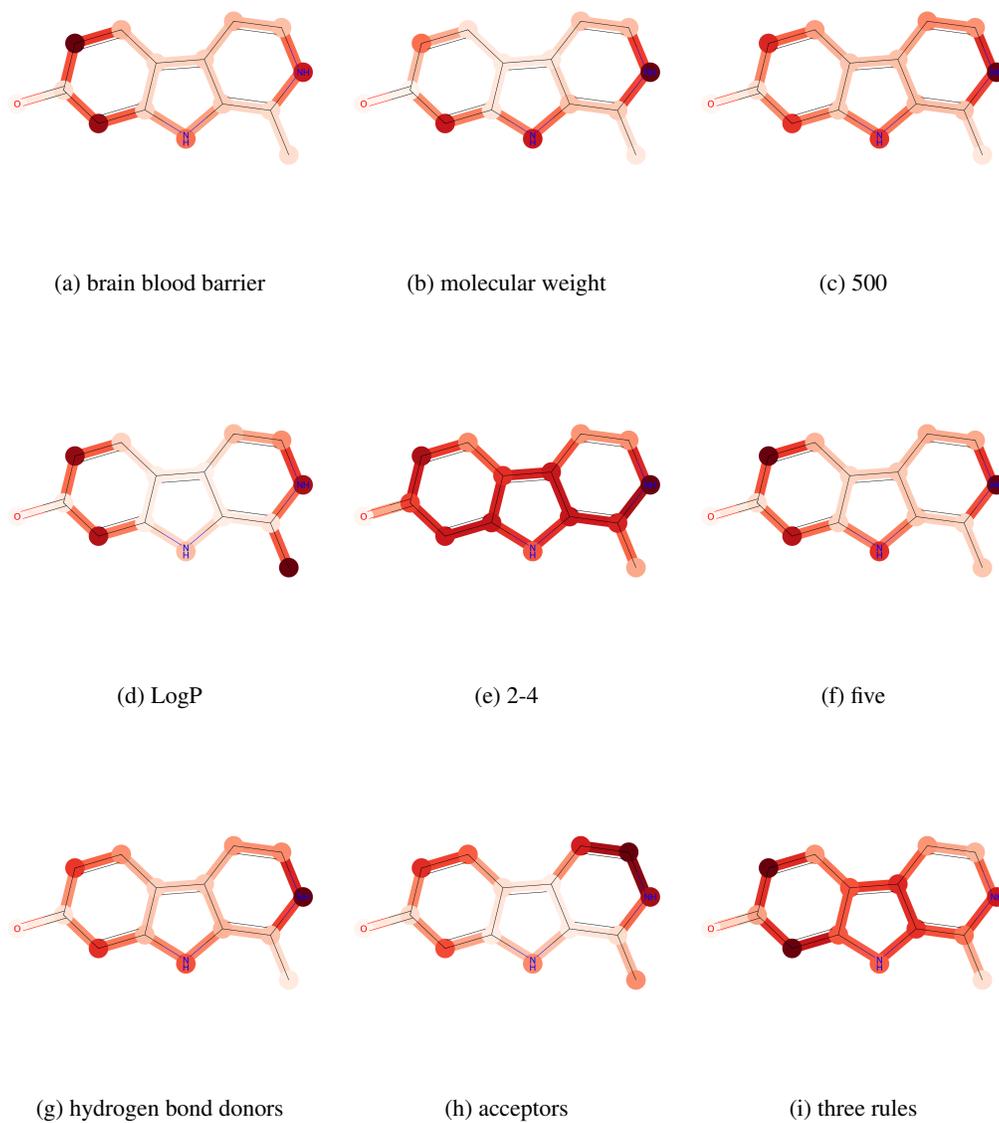


Figure 12: Visualization of attention for BBBP on molecule CC1=C2NC3=CC(=O)C=CC3=C2C=CN1

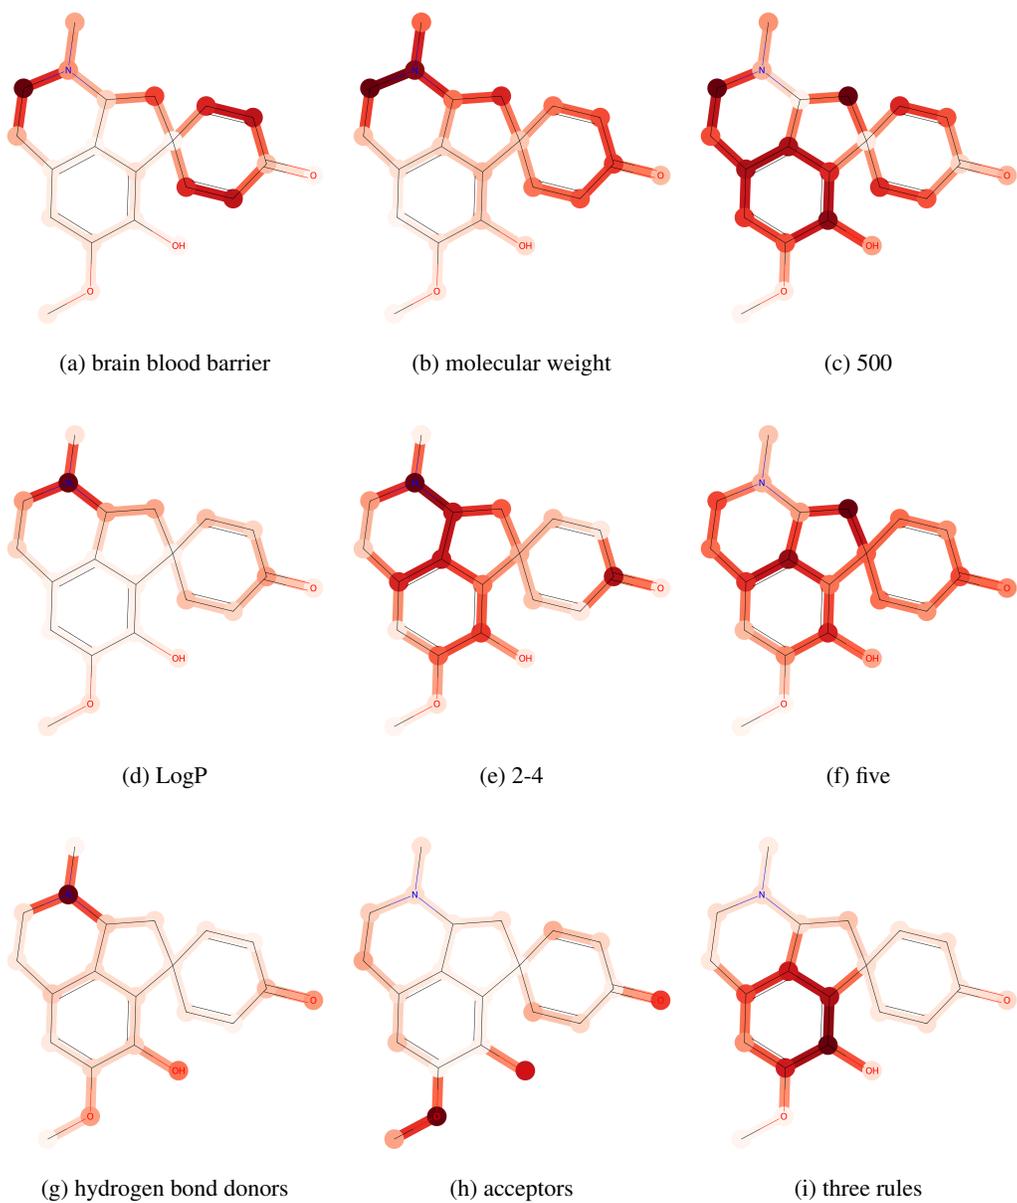


Figure 13: Visualization of attention for BBBP on molecule  
COc1cc2CCN(C)C3CC4(C=CC(=O)C=C4)c(c1O)c23

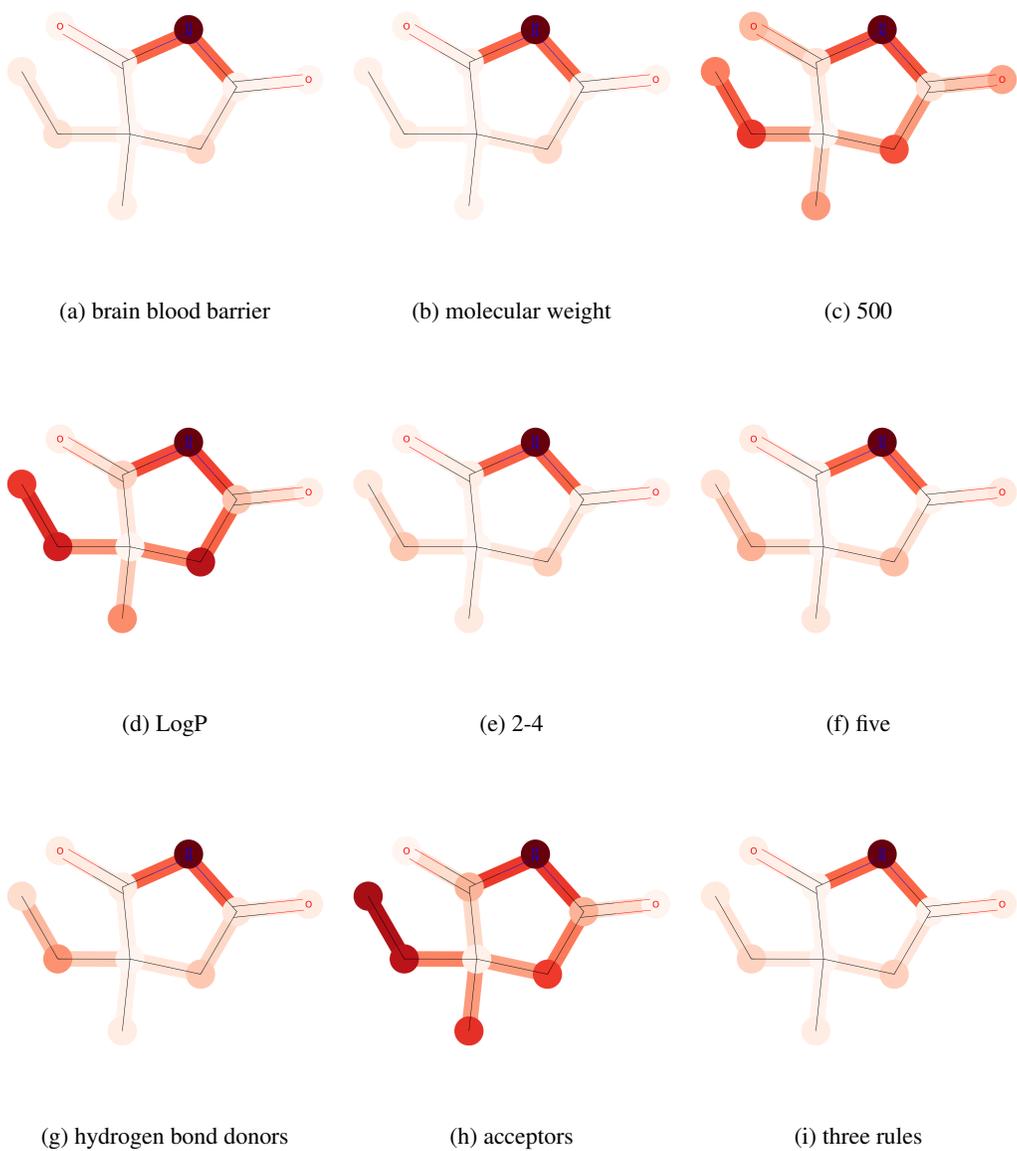


Figure 14: Visualization of attention for BBBP on molecule  
CCC1(C)CC(=O)NC1=O

## 520 References

- 521 [1] Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K.,  
522 Reynolds, M., et al. (2022). Flamingo: a visual language model for few-shot learning. *Advances in Neural*  
523 *Information Processing Systems*, 35:23716–23736.
- 524 [2] Bachlechner, T., Majumder, B. P., Mao, H., Cottrell, G., and McAuley, J. (2021). Rezero is all you need:  
525 Fast convergence at large depth. In *Uncertainty in Artificial Intelligence*, pages 1352–1361. PMLR.
- 526 [3] Bagal, V., Aggarwal, R., Vinod, P., and Priyakumar, U. D. (2021). Molgpt: molecular generation using a  
527 transformer-decoder model. *Journal of Chemical Information and Modeling*, 62(9):2064–2076.
- 528 [4] Bao, H., Dong, L., Piao, S., and Wei, F. (2021). Beit: Bert pre-training of image transformers. *arXiv preprint*  
529 *arXiv:2106.08254*.
- 530 [5] Bastos, A., Nadgeri, A., Singh, K., Kanezashi, H., Suzumura, T., and Mulang, I. O. (2022). Investigating  
531 expressiveness of transformer in spectral domain for graphs. *arXiv preprint arXiv:2201.09332*.
- 532 [6] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P.,  
533 Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information*  
534 *processing systems*, 33:1877–1901.
- 535 [7] Casanovas, R., Ortega-Castro, J., Frau, J., Donoso, J., and Munoz, F. (2014). Theoretical pka calculations  
536 with continuum model solvents, alternative protocols to thermodynamic cycles. *International Journal of*  
537 *Quantum Chemistry*, 114(20):1350–1363.
- 538 [8] Chen, D., O’Bray, L., and Borgwardt, K. (2022a). Structure-aware transformer for graph representation  
539 learning. In *International Conference on Machine Learning*, pages 3469–3489. PMLR.
- 540 [9] Chen, X., Wang, X., Changpinyo, S., Piergiovanni, A., Padlewski, P., Salz, D., Goodman, S., Grycner, A.,  
541 Mustafa, B., Beyer, L., et al. (2022b). Pali: A jointly-scaled multilingual language-image model. *arXiv*  
542 *preprint arXiv:2209.06794*.
- 543 [10] Chithrananda, S., Grand, G., and Ramsundar, B. (2020). Chemberta: Large-scale self-supervised pretraining  
544 for molecular property prediction. *arXiv preprint arXiv:2010.09885*.
- 545 [11] Choromanski, K., Lin, H., Chen, H., Zhang, T., Sehanobish, A., Likhoshesterov, V., Parker-Holder, J.,  
546 Sarlos, T., Weller, A., and Weingarten, T. (2022). From block-toeplitz matrices to differential equations on  
547 graphs: towards a general theory for scalable masked transformers. In *International Conference on Machine*  
548 *Learning*, pages 3962–3983. PMLR.
- 549 [12] Chung, H. W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, E., Wang, X., Dehghani, M., Brahma,  
550 S., et al. (2022). Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- 551 [13] Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional  
552 transformers for language understanding. In *Proceedings of the 2019 Conference of the North American*  
553 *Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long*  
554 *and Short Papers)*, pages 4171–4186.
- 555 [14] Dwivedi, V. P. and Bresson, X. (2020). A generalization of transformer networks to graphs. *arXiv preprint*  
556 *arXiv:2012.09699*.
- 557 [15] Edwards, C., Lai, T., Ros, K., Honke, G., and Ji, H. (2022). Translation between molecules and natural  
558 language. *arXiv preprint arXiv:2204.11817*.
- 559 [16] Edwards, C., Zhai, C., and Ji, H. (2021). Text2mol: Cross-modal molecule retrieval with natural language  
560 queries. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*,  
561 pages 595–607.
- 562 [17] Efrat, A. and Levy, O. (2020). The turking test: Can language models understand instructions? *arXiv*  
563 *preprint arXiv:2010.11982*.
- 564 [18] Fang, Y., Zhang, Q., Yang, H., Zhuang, X., Deng, S., Zhang, W., Qin, M., Chen, Z., Fan, X., and Chen, H.  
565 (2022). Molecular contrastive learning with chemical element knowledge graph. In *Proceedings of the AAAI*  
566 *Conference on Artificial Intelligence*, volume 36, pages 3968–3976.
- 567 [19] Gaulton, A., Bellis, L. J., Bento, A. P., Chambers, J., Davies, M., Hersey, A., Light, Y., McGlinchey, S.,  
568 Michalovich, D., Al-Lazikani, B., et al. (2012). ChEMBL: a large-scale bioactivity database for drug discovery.  
569 *Nucleic acids research*, 40(D1):D1100–D1107.

- 570 [20] Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for  
571 quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR.
- 572 [21] Gosselet, F., Loiola, R. A., Roig, A., Rosell, A., and Culot, M. (2021). Central nervous system delivery of  
573 molecules across the blood-brain barrier. *Neurochemistry International*, 144:104952.
- 574 [22] Guo, L., Zhang, Q., and Chen, H. (2022). Unleashing the power of transformer for graphs. *arXiv preprint*  
575 *arXiv:2202.10581*.
- 576 [23] Hassani, K. and Khasahmadi, A. H. (2020). Contrastive multi-view representation learning on graphs. In  
577 *International conference on machine learning*, pages 4116–4126. PMLR.
- 578 [24] Honda, S., Shi, S., and Ueda, H. R. (2019). Smiles transformer: Pre-trained molecular fingerprint for low  
579 data drug discovery. *arXiv preprint arXiv:1911.04738*.
- 580 [25] Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., and Leskovec, J. (2019). Strategies for  
581 pre-training graph neural networks. *arXiv preprint arXiv:1905.12265*.
- 582 [26] Kim, J., Nguyen, T. D., Min, S., Cho, S., Lee, M., Lee, H., and Hong, S. (2022). Pure transformers are  
583 powerful graph learners. *arXiv preprint arXiv:2207.02505*.
- 584 [27] Kreuzer, D., Beaini, D., Hamilton, W., Létourneau, V., and Tossou, P. (2021). Rethinking graph transformers  
585 with spectral attention. *Advances in Neural Information Processing Systems*, 34:21618–21629.
- 586 [28] Li, J., Li, D., Savarese, S., and Hoi, S. (2023). Blip-2: Bootstrapping language-image pre-training with  
587 frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.
- 588 [29] Liu, L., Liu, X., Gao, J., Chen, W., and Han, J. (2020). Understanding the difficulty of training transformers.  
589 In *2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020*, pages 5747–5763.  
590 Association for Computational Linguistics (ACL).
- 591 [30] Liu, S., Demirel, M. F., and Liang, Y. (2019a). N-gram graph: Simple unsupervised representation for  
592 graphs, with applications to molecules. *Advances in neural information processing systems*, 32.
- 593 [31] Liu, S., Nie, W., Wang, C., Lu, J., Qiao, Z., Liu, L., Tang, J., Xiao, C., and Anandkumar, A. (2022). Multi-  
594 modal molecule structure-text model for text-based retrieval and editing. *arXiv preprint arXiv:2212.10789*.
- 595 [32] Liu, S., Wang, H., Liu, W., Lasenby, J., Guo, H., and Tang, J. (2021). Pre-training molecular graph  
596 representation with 3d geometry. In *International Conference on Learning Representations*.
- 597 [33] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov,  
598 V. (2019b). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- 599 [34] Maziarka, Ł., Danel, T., Mucha, S., Rataj, K., Tabor, J., and Jastrzebski, S. (2020). Molecule attention  
600 transformer. *arXiv preprint arXiv:2002.08264*.
- 601 [35] Mialon, G., Chen, D., Selosse, M., and Mairal, J. (2021). Graphit: Encoding graph structure in transformers.  
602 *arXiv preprint arXiv:2106.05667*.
- 603 [36] Mishra, S., Khashabi, D., Baral, C., Choi, Y., and Hajishirzi, H. (2021). Reframing instructional prompts  
604 to gptk’s language. *arXiv preprint arXiv:2109.07830*.
- 605 [37] Mishra, S., Khashabi, D., Baral, C., and Hajishirzi, H. (2022). Cross-task generalization via natural  
606 language crowdsourcing instructions. In *Proceedings of the 60th Annual Meeting of the Association for*  
607 *Computational Linguistics (Volume 1: Long Papers)*, pages 3470–3487.
- 608 [38] Nielsen, M. H., Pedersen, F. S., and Kjems, J. (2005). Molecular strategies to inhibit hiv-1 replication.  
609 *Retrovirology*, 2:1–20.
- 610 [39] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama,  
611 K., Ray, A., et al. (2022). Training language models to follow instructions with human feedback. *Advances*  
612 *in Neural Information Processing Systems*, 35:27730–27744.
- 613 [40] Park, W., Chang, W.-G., Lee, D., Kim, J., et al. (2022). Grpe: Relative positional encoding for graph  
614 transformer. In *ICLR2022 Machine Learning for Drug Discovery*.
- 615 [41] Parmar, M., Mishra, S., Purohit, M., Luo, M., Mohammad, M., and Baral, C. (2022). In-boxbart: Get  
616 instructions into biomedical multi-task learning. In *Findings of the Association for Computational Linguistics:*  
617 *NAACL 2022*, pages 112–128.

- 618 [42] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2020).  
619 Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine*  
620 *Learning Research*, 21(1):5485–5551.
- 621 [43] Ramsundar, B., Eastman, P., Walters, P., and Pande, V. (2019). *Deep learning for the life sciences: applying*  
622 *deep learning to genomics, microscopy, drug discovery, and more*. O’Reilly Media.
- 623 [44] Rong, Y., Bian, Y., Xu, T., Xie, W., Wei, Y., Huang, W., and Huang, J. (2020). Self-supervised graph  
624 transformer on large-scale molecular data. *Advances in Neural Information Processing Systems*, 33:12559–  
625 12571.
- 626 [45] Ross, J., Belgodere, B., Chenthamarakshan, V., Padhi, I., Mroueh, Y., and Das, P. (2022). Molformer:  
627 Large scale chemical language representations capture molecular structure and properties.
- 628 [46] Sanh, V., Webson, A., Raffel, C., Bach, S. H., Sutawika, L., Alyafeai, Z., Chaffin, A., Stiegler, A., Scio,  
629 T. L., Raja, A., et al. (2021). Multitask prompted training enables zero-shot task generalization. *arXiv preprint*  
630 *arXiv:2110.08207*.
- 631 [47] Schick, T. and Schütze, H. (2021). Exploiting cloze-questions for few-shot text classification and natural  
632 language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for*  
633 *Computational Linguistics: Main Volume*, pages 255–269.
- 634 [48] Seidl, P., Vall, A., Hochreiter, S., and Klambauer, G. (2023). Enhancing activity prediction models in drug  
635 discovery with the ability to understand human language. *arXiv preprint arXiv:2303.03363*.
- 636 [49] Su, B., Du, D., Yang, Z., Zhou, Y., Li, J., Rao, A., Sun, H., Lu, Z., and Wen, J.-R. (2022). A molec-  
637 ular multimodal foundation model associating molecule graphs with natural language. *arXiv preprint*  
638 *arXiv:2209.05481*.
- 639 [50] Sun, F.-Y., Hoffman, J., Verma, V., and Tang, J. (2020). Infograph: Unsupervised and semi-supervised  
640 graph-level representation learning via mutual information maximization. In *International Conference on*  
641 *Learning Representations*.
- 642 [51] Sun, M., Xing, J., Wang, H., Chen, B., and Zhou, J. (2021). Mocl: Data-driven molecular fingerprint via  
643 knowledge-aware contrastive learning from molecular graph. In *Proceedings of the 27th ACM SIGKDD*  
644 *Conference on Knowledge Discovery & Data Mining*, pages 3585–3594.
- 645 [52] Sun, R., Dai, H., and Yu, A. W. (2022). Does gnn pretraining help molecular representation? *Advances in*  
646 *Neural Information Processing Systems*, 35:12096–12109.
- 647 [53] Suresh, S., Li, P., Hao, C., and Neville, J. (2021). Adversarial graph augmentation to improve graph  
648 contrastive learning. *Advances in Neural Information Processing Systems*, 34:15920–15933.
- 649 [54] Taylor, R., Kardas, M., Cucurull, G., Scialom, T., Hartshorn, A., Saravia, E., Poulton, A., Kerkez, V., and  
650 Stojnic, R. (2022). Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085*.
- 651 [55] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I.  
652 (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- 653 [56] Velickovic, P., Fedus, W., Hamilton, W. L., Liò, P., Bengio, Y., and Hjelm, R. D. (2019). Deep graph  
654 infomax. *ICLR (Poster)*, 2(3):4.
- 655 [57] Wang, S., Guo, Y., Wang, Y., Sun, H., and Huang, J. (2019). Smiles-bert: large scale unsupervised  
656 pre-training for molecular property prediction. In *Proceedings of the 10th ACM international conference on*  
657 *bioinformatics, computational biology and health informatics*, pages 429–436.
- 658 [58] Wang, Y., Magar, R., Liang, C., and Barati Farimani, A. (2022a). Improving molecular contrastive learning  
659 via faulty negative mitigation and decomposed fragment contrast. *Journal of Chemical Information and*  
660 *Modeling*, 62(11):2713–2725.
- 661 [59] Wang, Y., Min, Y., Shao, E., and Wu, J. (2021). Molecular graph contrastive learning with parameterized  
662 explainable augmentations. In *2021 IEEE International Conference on Bioinformatics and Biomedicine*  
663 *(BIBM)*, pages 1558–1563. IEEE.
- 664 [60] Wang, Y., Mishra, S., Alipoormolabashi, P., Kordi, Y., Mirzaei, A., Naik, A., Ashok, A., Dhanasekaran,  
665 A. S., Arunkumar, A., Stap, D., et al. (2022b). Super-naturalinstructions: Generalization via declarative  
666 instructions on 1600+ nlp tasks. In *Proceedings of the 2022 Conference on Empirical Methods in Natural*  
667 *Language Processing*, pages 5085–5109.

- 668 [61] Withnall, M., Chen, H., and Tetko, I. V. (2018). Matched molecular pair analysis on large melting point  
669 datasets: a big data perspective. *ChemMedChem*, 13(6):599–606.
- 670 [62] Wu, Z., Jain, P., Wright, M., Mirhoseini, A., Gonzalez, J. E., and Stoica, I. (2021). Representing long-range  
671 context for graph neural networks with global attention. *Advances in Neural Information Processing Systems*,  
672 34:13266–13279.
- 673 [63] Xia, J., Wu, L., Chen, J., Hu, B., and Li, S. Z. (2022). Simgrace: A simple framework for graph contrastive  
674 learning without data augmentation. In *Proceedings of the ACM Web Conference 2022*, pages 1070–1079.
- 675 [64] Xu, D., Cheng, W., Luo, D., Chen, H., and Zhang, X. (2021). Infogcl: Information-aware graph contrastive  
676 learning. *Advances in Neural Information Processing Systems*, 34:30414–30425.
- 677 [65] Ying, C., Cai, T., Luo, S., Zheng, S., Ke, G., He, D., Shen, Y., and Liu, T.-Y. (2021). Do transformers really  
678 perform badly for graph representation? *Advances in Neural Information Processing Systems*, 34:28877–  
679 28888.
- 680 [66] You, Y., Chen, T., Shen, Y., and Wang, Z. (2021). Graph contrastive learning automated. In *International  
681 Conference on Machine Learning*, pages 12121–12132. PMLR.
- 682 [67] You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., and Shen, Y. (2020). Graph contrastive learning with  
683 augmentations. *Advances in neural information processing systems*, 33:5812–5823.
- 684 [68] Zeng, Z., Yao, Y., Liu, Z., and Sun, M. (2022). A deep-learning system bridging molecule structure and  
685 biomedical text with comprehension comparable to human professionals. *Nature communications*, 13(1):862.
- 686 [69] Zhang, J., Zhang, H., Xia, C., and Sun, L. (2020). Graph-bert: Only attention is needed for learning graph  
687 representations. *arXiv preprint arXiv:2001.05140*.
- 688 [70] Zhang, Z., Liu, Q., Wang, H., Lu, C., and Lee, C.-K. (2021). Motif-based graph self-supervised learning  
689 for molecular property prediction. *Advances in Neural Information Processing Systems*, 34:15870–15882.
- 690 [71] Zhao, H., Ma, S., Zhang, D., Deng, Z.-H., and Wei, F. (2022). Are more layers beneficial to graph  
691 transformers? In *The Eleventh International Conference on Learning Representations*.
- 692 [72] Zhong, R., Lee, K., Zhang, Z., and Klein, D. (2021). Adapting language models for zero-shot learning by  
693 meta-tuning on dataset and prompt collections. In *Findings of the Association for Computational Linguistics:  
694 EMNLP 2021*, pages 2856–2878.