

Learning Multi-view Molecular Representations with Structured and Unstructured Knowledge

Yizhen Luo
Institute of AI Industry Research
(AIR), Tsinghua University
Beijing, China
yz-luo22@mails.tsinghua.edu.cn

Kai Yang
Institute of AI Industry Research
(AIR), Tsinghua University
Beijing, China
yangkai@air.tsinghua.edu.cn

Massimo Hong
Institute of AI Industry Research
(AIR), Tsinghua University
Beijing, China
hongcd21@mails.tsinghua.edu.cn

Xing Yi Liu
Institute of AI Industry Research
(AIR), Tsinghua University
Beijing, China
liuxingyi99@gmail.com

Zikun Nie
Institute of AI Industry Research
(AIR), Tsinghua University
Beijing, China
nzk20@mails.tsinghua.edu.cn

Hao Zhou
Institute of AI Industry Research
(AIR), Tsinghua University
Beijing, China
zhouhao@air.tsinghua.edu.cn

Zaiqing Nie*
Institute of AI Industry Research
(AIR), Tsinghua University
Pharmolix Inc.
Beijing, China
zaiqing@air.tsinghua.edu.cn

ABSTRACT

Capturing molecular knowledge with representation learning approaches holds significant potential in vast scientific fields such as chemistry and life science. An effective and generalizable molecular representation is expected to capture the consensus and complementary molecular expertise from diverse views and perspectives. However, existing works fall short in learning multi-view molecular representations, due to challenges in explicitly incorporating view information and handling molecular knowledge from heterogeneous sources. To address these issues, we present MV-Mol, a molecular representation learning model that harvests multi-view molecular expertise from chemical structures, unstructured knowledge from biomedical texts, and structured knowledge from knowledge graphs. We utilize text prompts to model view information and design a fusion architecture to extract view-based molecular representations. We develop a two-stage pre-training procedure, exploiting heterogeneous data of varying quality and quantity. Through extensive experiments, we show that MV-Mol provides improved representations that substantially benefit molecular property prediction. Additionally, MV-Mol exhibits state-of-the-art performance in multi-modal comprehension of molecular structures and texts. Code and data are available at <https://github.com/PharMolix/OpenBioMed>.

*Corresponding author.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
KDD '24, August 25–29, 2024, Barcelona, Spain.
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0490-1/24/08
<https://doi.org/10.1145/3637528.3672043>

CCS CONCEPTS

• **Applied computing** → **Bioinformatics**; • **Computing methodologies** → *Knowledge representation and reasoning*; *Natural language processing*.

KEYWORDS

Multi-view Molecular Representation Learning, Knowledge Graphs, Text Mining

ACM Reference Format:

Yizhen Luo, Kai Yang, Massimo Hong, Xing Yi Liu, Zikun Nie, Hao Zhou, and Zaiqing Nie. 2024. Learning Multi-view Molecular Representations with Structured and Unstructured Knowledge. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3672043>

1 INTRODUCTION

Understanding the properties and functions of small molecules is a pivotal issue in various scientific applications such as chemistry [2], biology [56], and material design [15]. Due to the substantial annotation costs for molecules, molecular representation learning (MRL), which aims at capturing molecular knowledge from abundant unlabeled data, has attracted significant research attention. Through self-supervised learning [33] on various types of molecular structures, such as 1D SMILES strings [5, 18], 2D graphs [17, 67, 73] or 3D conformations [57, 80, 81], MRL models have achieved great success in advancing our understanding of molecules.

However, to further extend the application scope of MRL models, one is faced with a critical problem: molecular expertise is complicated and multifaceted, spanning diverse disciplines and views. It comprises consensus information shared by multiple views, as well as complementary information within each specific view [3]. As

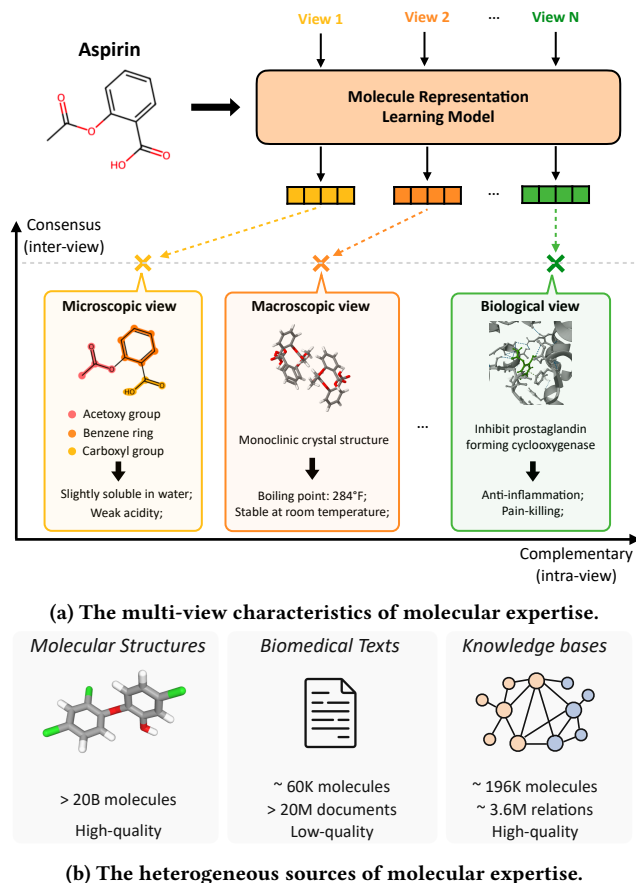


Figure 1: An overview of molecular expertise. Molecular expertise covers diverse views that share consensus and complementary information. It resides within heterogeneous sources that vary in quality and quantity.

exemplified in Figure 1a, the characteristics of Aspirin are identified through various perspectives. From a microscopic view, the molecule comprises a benzene ring, an acetoxy group, and a carboxyl group, resulting in slight water solubility and weak acidity. From a macroscopic view, it exhibits a monoclinic crystal structure, influencing physical properties such as boiling point and stability. From a biological view, it inhibits the production of prostaglandins by targeting cyclooxygenase, thus exhibiting anti-inflammatory and pain-killing functions [63]. A versatile MRL model is expected to generate view-based representations to address the distinctions between different application contexts. Unfortunately, existing works primarily focus on learning a universal representation to capture the consensus information across various views, and fall short in grasping the complementary characteristics of each view.

More recently, the emergence of MRL methods incorporating molecular knowledge from heterogeneous sources has opened novel avenues for addressing the multi-view problem. These models focus on jointly comprehending molecular structures, structured knowledge from knowledge bases, and unstructured knowledge from biomedical texts. The integration of heterogeneous inputs is accomplished either through language modeling objectives within

a unified model [36, 49, 74], or by contrastive learning objectives with independent encoders [31, 55, 58]. Innovated by the success of these models, we aim to capture multi-view molecular expertise from these structured and unstructured knowledge sources.

Nevertheless, there exist two challenges in learning multi-view molecular representations. First, as illustrated in Figure 1a, the MRL model should incorporate view information explicitly into its representations to ensure its adaptability in broad applications. However, prior MRL models integrate view information implicitly with ‘wrapped texts’ [36, 49] or fine-tuning on downstream tasks, which compromises their understanding of the consensus and complementary relationships between molecular knowledge from different views. Besides, as shown in Figure 1b, it is essential to address the heterogeneity of information sources, including molecular structures, biomedical texts, and knowledge graphs, which vary in quality and quantity. Previous works treat structured and unstructured knowledge indiscriminately by transforming knowledge graphs into texts. However, this may introduce biases across different views due to the imbalanced distribution of pre-training data.

In this work, we propose MV-Mol, a comprehensive framework for **Multi-View Molecular representation learning** with structured and unstructured knowledge, to address the aforementioned problems. To explicitly incorporate view information, we utilize text prompts to capture the complementary and consensus characteristics of different views for molecules. We leverage Q-Former [25], a multi-modal fusion architecture, to extract view-based molecular representations by jointly comprehending molecular structures and view prompts. Then, we propose a two-stage pre-training strategy to address the heterogeneity of structured and unstructured knowledge. The first stage aligns molecular structures with large-scale, noisy texts, extracting consensus information across comprehensive views. The second stage incorporates high-quality, structured knowledge from knowledge graphs. Drawing inspiration from knowledge-enhanced pre-training [65, 66], we treat relations as specific types of views and describe them with texts. With contrastive and generative objectives, MV-Mol is endowed with the ability to grasp complementary information within different views.

We show the superior performance of MV-Mol through fine-tuning across various downstream tasks. Benefiting from view-based molecular representations and two-stage pre-training, MV-Mol achieves an average of 1.24% absolute gains over the state-of-the-art method Uni-Mol [80] on molecule property prediction [70]. Additionally, MV-Mol exhibits a deeper understanding of connections between molecular structures and texts. On cross-modal retrieval, MV-Mol improves the top-1 retrieval accuracy by 12.9% on average over the best-performing baselines. On cross-modal generation [10], MV-Mol also yields more accurate predictions, as validated by qualitative and quantitative studies.

Our contributions are summarized as follows: (1) To the best of our knowledge, MV-Mol is the first work that addresses the multi-view problem in molecular representation learning. (2) We propose to incorporate view information explicitly by jointly comprehending molecular structures and view prompts within a fusion network. (3) We develop a two-stage training pipeline to harvest multi-view molecular expertise from structured and unstructured knowledge sources that vary in quality and quantity.

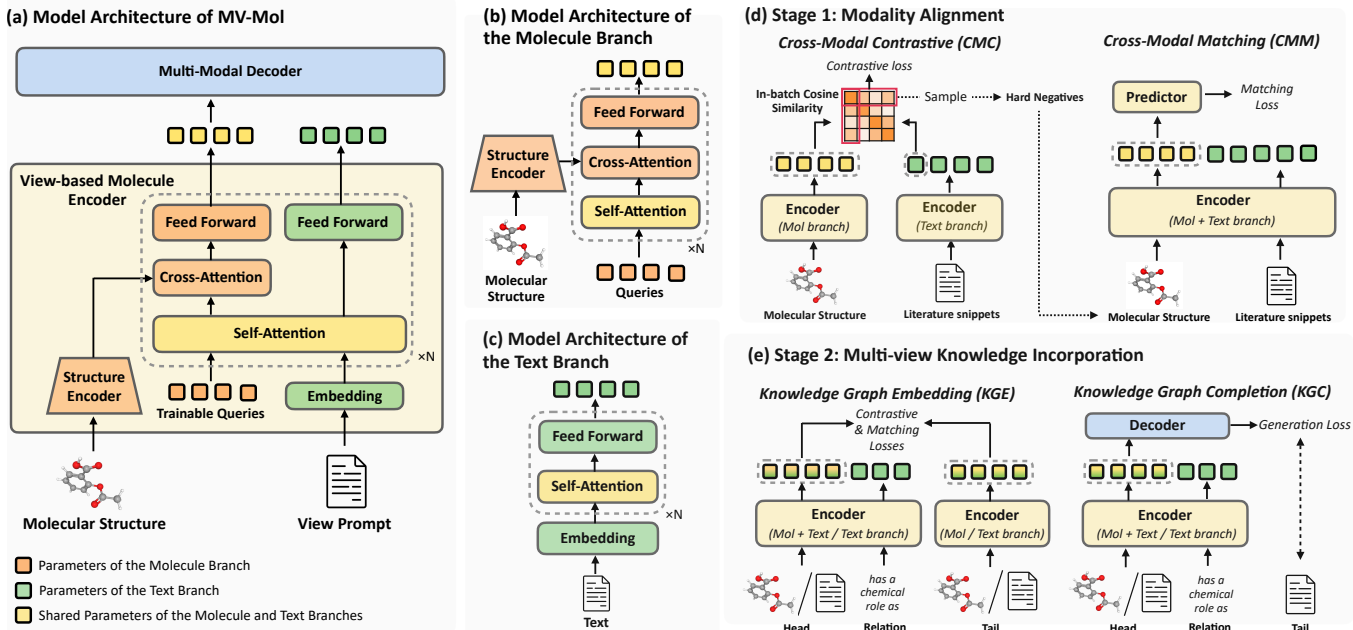


Figure 2: Model architecture and pre-training pipeline of MV-Mol. (a) MV-Mol is composed of a view-based molecule encoder and a multi-modal decoder. (b) The molecule branch of the view-based molecule encoder. (c) The text branch of the view-based molecule encoder. (d) We perform cross-modal contrastive and cross-modal matching for modality alignment. (e) We model relation as a textual prompt that constrains molecular knowledge from a specific view, and design knowledge graph embedding and knowledge graph completion objectives for multi-view knowledge incorporation. Both branches are activated when the head entity refers to a molecule.

2 RELATED WORK

2.1 Multi-view Molecular Representation Learning

Multi-view representation learning [29, 79] aims to pursue a unified and comprehensive feature representation based on multi-view data. Most MRL models treat different forms of molecular structures as distinct views and model complementary and consensus information with different pre-training objectives. For example, DVMP [82] treats 1D SMILES sequences and 2D molecular graphs as dual views. GraphMVP [32] treats 2D topologies and 3D geometries as two views, and proposes independent objectives for intra-view and inter-view modeling. MORN [41] treats the sequential, topological, and spatial information as different views and implements a two-stage adaptive learning strategy with multi-view fusion for molecular property prediction. [64] treats atoms, scaffolds, and functional groups as independent views and proposes a prompt-based aggregation module to extract task-specific molecular representations. In this work, we extend the multi-view concept of molecules as their inherent properties and patterns within a certain context, which is defined flexibly by human-understandable textual prompts.

2.2 Molecular Representation Learning with Heterogeneous Sources

Recently, MRL approaches that jointly harvest heterogeneous sources including molecular structures, biomedical texts, and knowledge graphs have surged. Drawing parallels from vision-language pre-training (VLP) [25, 51, 59], these models can be categorized as

follows: (1) Generative models including KV-PLM [74], MolT5 [10], MolXPT [36], Text+Chem [6], GIMLET [77], BioMedGPT [40], Mol-Instructions [13], BioT5 [49] and BioT5+ [48]. These models treat molecules and texts as two languages and jointly process them within a unified language model [21, 52, 53]. (2) Contrastive models like MoMu [58], MoleculeSTM [31], CLAMP [55], KANO [14] and GODE [19]. These models incorporate texts or knowledge graphs as cross-modal supervision signals and augment molecular representations with contrastive learning [51]. (3) Hybrid models including MolFM [39], MolCA [35], GIT-Mol [30] and 3D-MoLM [28]. These models perform fine-grained feature integration with multi-modal fusion networks and incorporate both generative and contrastive objectives for pre-training. While MV-Mol also adopts a fusion network architecture, we distinguish it from prior works by our explicit modeling of views by jointly comprehending molecular structures and texts and our pre-training process that addresses data imbalance of heterogeneous knowledge sources.

3 METHOD

As depicted in Figure 2, MV-Mol is designed to learn multi-view molecular representations from heterogeneous sources. The core of our approach lies in modeling view information with text prompts.

The remainder of this section is organized as follows: First, we provide a roadmap for extracting multi-view molecular expertise within heterogeneous sources from a knowledge-enhanced pre-training perspective (Sec. 3.1). Then, we introduce the model architecture for MV-Mol (Sec. 3.2) that extracts view-based molecular representations. To address the heterogeneity of pre-training data,

we develop a two-stage pre-training paradigm, harvesting unstructured knowledge from biomedical documents through modality alignment (Sec. 3.3) and structured knowledge from knowledge bases with knowledge incorporation (Sec. 3.4).

3.1 A Knowledge-enhanced Pre-training Perspective for Multi-view MRL

In this section, we interpret multi-view MRL from the perspective of knowledge-enhanced pre-training [16, 65, 66], providing a roadmap for methodology design. Inspired by OntoProtein [76], we formulate the pre-training data as triplets (h, r, t) , where h and t are head and tail entities, and r signifies the relation. We define two types of entities, E_{mol} and E_{text} . $e^{(s)} \in E_{\text{mol}}$ is denoted as a molecule node with structural information. $e^{(t)} \in E_{\text{text}}$ is denoted as a text node formulated by a token sequence. r is also formulated as text tokens that describe the relation. Depending on the types of the head and tail entities, we further categorize pre-training data as molecule-text triplets, molecule-molecule triplets, and text-text triplets.

Knowledge-enhanced pre-training methods typically maximize the following objective:

$$\mathcal{L} = S(f(h, r), g(t)), \quad (1)$$

where $f(h, r)$ is the relation-transformed representation of the head entity, $g(t)$ is the representation of the tail entity, and $S(\cdot, \cdot)$ is a scoring function that measures the alignment of the triplet. We observe for $h \in E_{\text{mol}}$, the view-based molecular representations have an equivalent form of $f(h, r)$. Besides, structured knowledge graphs are indispensable sources for multi-view MRL, as both views and relations can be conceptualized as textual refinements over specific aspects of molecules. Therefore, multi-view MRL comprises the following key design components:

- The model architecture for $f(h, r)$, tasked with extracting different feature representations of h in the context of r . As h and r may come from different modalities, the model should jointly comprehend molecular structures and texts.
- The formulation of the scoring function $S(\cdot, \cdot)$, intended to harvest rich supervision signals from heterogeneous sources.

3.2 Model Architecture for MV-Mol

The model architecture of MV-Mol is depicted in Figure 2(a), comprising 2 parts: (1) A view-based molecule encoder that jointly comprehends molecular structures and texts. (2) A multi-modal decoder that generates human-understandable texts.

View-based Molecule Encoder. We define the structure of a molecule as $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{C})$ where \mathcal{V} represents atoms, \mathcal{E} represents bonds, and $\mathcal{C} \in \mathbb{R}^{|\mathcal{V}| \times 3}$ represents the 3D coordinates for each atom. We leverage Uni-Mol [80], a 15-layer Transformer [71] pre-trained on 209M 3D conformations, to encode \mathcal{M} . The pre-trained molecular model mitigates the disproportionality between molecular structures and text data [58]. The structure encoder f_{mol} transforms \mathcal{M} into feature representations for each atom, which we denote as $z^{(a)}$.

In MV-Mol, we use text prompts to model different views of molecules, represented by a sequence of L tokens $\mathcal{T} = [x_1, x_2, \dots, x_L]$. Consequently, the view-based molecule encoder is expected to comprehend molecular structures and texts simultaneously. To achieve

this, we leverage Querying Transformer (Q-Former) [25], a novel multi-modal fusion architecture from vision-language pre-training. The Q-Former consists of two branches, each consisting of 12 bi-directional Transformer layers initialized with KV-PLM* [74]. The molecule branch takes K trainable query vectors as input embeddings. The query features extract pivotal information from atomic representations $z^{(a)}$ through cross-attention within every other Transformer layer. The text branch digests \mathcal{T} as its input. The self-attention layers across the two branches are shared, allowing the queries to grasp fine-grained semantics from text prompts. In this way, the molecular representations are propagated within each Q-Former layer to incorporate information from different views.

Overall, the view-based molecule encoder f_v can transform the molecular structure to a fixed number of features with the structure encoder and the molecule branch of Q-Former, which is illustrated in Fig. 2b:

$$z^{(s)} = [z_1^{(s)}, z_2^{(s)}, \dots, z_K^{(s)}] = f_v(\mathcal{M}). \quad (2)$$

Using Q-Former’s text branch, it can also comprehend natural language as in Fig. 2c:

$$z^{(t)} = [z_1^{(t)}, z_2^{(t)}, \dots, z_L^{(t)}] = f_v(\mathcal{T}), \quad (3)$$

Most importantly, as shown in Figure 2a, it can grasp view-based molecular representations through multi-modal feature fusion of molecular structures and texts:

$$z^{(s,t)} = [z_1^{(s,t)}, z_2^{(s,t)}, \dots, z_K^{(s,t)}] = f_v(\mathcal{M}, \mathcal{T}). \quad (4)$$

Multi-modal Decoder. The multi-modal decoder allows MV-Mol to interpret the view-based molecular representations with natural language. We employ BioT5 [49], a molecular language model with 12 Transformer encoder layers and 12 Transformer decoder layers. The model has undergone multi-task pre-training on SELFIES strings of molecular structures, FASTA sequences of proteins, and biomolecule-related texts. For MV-Mol, we leverage the decoder branch of BioT5 as the text decoder f_{dec} , which transforms the Q-Former outputs into a sequence of N tokens $y = [y_1, y_2, \dots, y_N]$ by casual generation:

$$y_i = \arg\max \left\{ P \left(y_i | y_{<i}, z^{(\cdot)} \right) \right\}, \quad (5)$$

where $y_{<i}$ denotes preceding tokens, $z^{(\cdot)}$ refers to either $z^{(s)}$, $z^{(t)}$ or $z^{(s,t)}$, and the conditional probability is modeled by the text decoder f_{dec} .

Discussions. While our model shares a similar architecture with MolCA [35] and 3D-MoLM [28], we emphasize that MV-Mol extracts view-based molecular representations in Eq. 4 with both branches of the Q-Former. In contrast, prior works are unaware of view information and encode molecules as in Eq. 2, where the text branch of the Q-Former is deactivated.

3.3 Modality Alignment with Large-scale Unstructured Knowledge

In this stage, we aim to capture the consensus information within numerous views of molecules and cultivate MV-Mol’s capability of jointly comprehending molecular structures and texts. This is accomplished by aligning the feature space between molecular structures $h \in E_{\text{mol}}$ and the relevant contexts $t \in E_{\text{text}}$ from biomedical

literature. We omit the relation r due to the unaffordable costs of annotating millions of noisy texts. As shown in Figure 2b, the pre-training objectives are:

Cross-Modal Contrastive Loss. This objective maximizes the mutual information between the structural and textual representations. We define structure-text similarity as follows:

$$\text{sim}(z^{(s)}, z^{(t)}) = \max_{i=1,2,\dots,K} \left\{ f_{\text{proj}}(z_i^{(s)})^T f_{\text{proj}}(z_1^{(t)}) \right\}, \quad (6)$$

where $f_{\text{proj}}(\cdot)$ is a projection network composed of a fully-connected layer and l_2 normalization, and $z_1^{(t)}$ denotes the textual representation of the [CLS] token.

The objective function is calculated as:

$$\mathcal{L}_{\text{cmc}} = -\frac{1}{2} \left[\log \frac{\exp(\text{sim}[f_v(h), f_v(t)]/\tau)}{\sum_{h'} \exp(\text{sim}[f_v(h'), f_v(t)]/\tau)} + \log \frac{\exp(\text{sim}[f_v(h), f_v(t)]/\tau)}{\sum_{t'} \exp(\text{sim}[f_v(h), f_v(t')]/\tau)} \right], \quad (7)$$

where h' and t' are molecules and texts within the same mini-batch, and τ is a temperature hyper-parameter.

Cross-Modal Matching Loss. This objective fosters the fine-grained comprehension of molecular structures and texts by predicting whether they correspond to the same molecule. Following [25, 27], we acquire hard negatives for each sample by obtaining one misaligned structure for the text and one misaligned text for the structure from the same mini-batch based on the similarity scores defined in Eq. 6. The objective is calculated as follows:

$$\mathcal{L}_{\text{cmm}} = H[y_{\text{cmm}}(h, t), f_{\text{cmm}}(f_v(h, t))], \quad (8)$$

where f_{cmm} is a predictor composed of an average pooling and a fully-connected layer, $y_{\text{cmm}}(\cdot, \cdot)$ is the ground truth label indicating the alignment of the molecule-text pair, and $H(\cdot, \cdot)$ denotes cross-entropy.

The overall objective is the summation of the two loss functions:

$$\mathcal{L}_1 = \mathbb{E}_{(h,t) \sim \mathcal{D}_1} (\mathcal{L}_{\text{cmc}} + \mathcal{L}_{\text{cmm}}), \quad (9)$$

where \mathcal{D}_1 is the pre-training dataset.

3.4 Multi-view Knowledge Incorporation with Heterogeneous Triplets

In this stage, we aim to incorporate high-quality structured knowledge into MV-Mol and capture the multi-view characteristics of molecules within three different types of triplets. Motivated by [65, 66], we introduce relation transformed representations of the head entity, $z^{(h,r)}$, calculated as follows:

$$z^{(h,r)} = \begin{cases} f_v(h, r), & h \in E_{\text{mol}} \\ f_v(h \oplus r), & h \in E_{\text{text}} \end{cases} \quad (10)$$

where \oplus denotes concatenation. By modeling relations as view prompts, we capture view-based molecular representations within the first term of the equation. We connect $z^{(h,r)}$ and t with objectives shown in Figure 2c, which we detail as follows:

Knowledge Graph Embedding Loss. This objective captures the global features within knowledge bases by aligning $z^{(h,r)}$ and the

representations of the tail entity. We first employ the following contrastive objective:

$$\mathcal{L}_{\text{kge_c}} = -\frac{1}{2} \left[\log \frac{\exp(\text{sim}_{\text{tri}}[z^{(h,r)}, f_v(t)]/\tau)}{\sum_{h'} \exp(\text{sim}_{\text{tri}}[z^{(h',r)}, f_v(t)]/\tau)} + \log \frac{\exp(\text{sim}_{\text{tri}}[z^{(h,r)}, f_v(t)]/\tau)}{\sum_{t'} \exp(\text{sim}_{\text{tri}}[z^{(h,r)}, f_v(t')]/\tau)} \right], \quad (11)$$

where h' and t' are head and tail entities within the same batch, and $\text{sim}_{\text{tri}}(\cdot, \cdot)$ calculates the representation similarity, which we detail in Appendix A. This objective is a generalized form of Eq. 7 that incorporates relation information from heterogeneous triplets.

Then, we incorporate a matching loss similar to Eq. 8 for molecule-text and text-text triplets with the same negative sampling strategy for the head and tail entity. We omit molecule-molecule triplets as f_v cannot process two molecules simultaneously. The objective is calculated as follows:

$$\mathcal{L}_{\text{kge_m}} = H[y_{\text{cmm}}(h, r, t), f_{\text{cmm}}(f_v(h, r \oplus t))]. \quad (12)$$

Knowledge Graph Completion Loss. This objective elicits MV-Mol to 'talk' about molecular knowledge from a specific view. We implement the language model-based tail entity prediction task [34, 50] by feeding the relation-transformed representations $z^{(h,r)}$ into the text decoder to generate descriptions for the tail entity. The objective is calculated for molecule-text and text-text triplets in an auto-regressive manner:

$$\mathcal{L}_{\text{kge}} = \sum_{x_i \in t} H[x_i, P(x_i | x_{<i}, z^{(h,r)})], \quad (13)$$

where the conditional probability is modeled by the text decoder f_{dec} .

The overall objective is calculated as follows:

$$\mathcal{L}_2 = \mathbb{E}_{(h,r,t) \in \mathcal{D}_2} (\mathcal{L}_{\text{kge_c}} + \mathcal{L}_{\text{kge_m}} + \mathcal{L}_{\text{kge}}), \quad (14)$$

where \mathcal{D}_2 is the knowledge graph.

4 EXPERIMENTS

In this section, we conduct extensive experiments to demonstrate the effectiveness of the proposed model by answering the following questions:

- **Q1.** Compared to prior MRL models, does MV-Mol provide better molecular representations?
- **Q2.** Does MV-Mol address the heterogeneity of molecular structures and texts?
- **Q3.** Do view-based molecular representations capture the consensus and complementary knowledge from different views? How do view prompts affect the representations?
- **Q4.** How do the design components within the two pre-training stages affect MV-Mol?

4.1 Pre-training Setup

Pre-training Dataset. For the modality alignment stage, we collect large-scale molecule-text pairs by named entity recognition (NER) [43] and entity linking [23] on 3.5M scientific publications following previous works [36, 49, 74], and obtain a total of 60K molecules and 12M molecule-text pairs. For the knowledge incorporation stage, we construct a knowledge graph by combining public databases

Table 1: AUROC scores for molecular property prediction on MoleculeNet. The best results are marked in bold, and the second-best results are underlined. \uparrow : the higher the better. \dagger : Our implementation. -: Not reported in the original paper. w/o view: without view prompts.

Model	BBBP \uparrow	Tox21 \uparrow	ToxCast \uparrow	SIDER \uparrow	ClinTox \uparrow	MUV \uparrow	HIV \uparrow	BACE \uparrow	Avg \uparrow
GraphCL [73]	67.5 \pm 3.3	75.0 \pm 0.3	62.8 \pm 0.2	60.1 \pm 1.3	78.9 \pm 4.2	77.1 \pm 1.0	75.0 \pm 0.4	68.7 \pm 7.8	70.64
GraphMVP [32]	72.4 \pm 1.6	74.4 \pm 0.2	63.1 \pm 0.4	63.9 \pm 1.2	77.5 \pm 4.2	75.0 \pm 1.0	77.0 \pm 1.2	81.2 \pm 0.9	73.07
GEM [12]	72.4 \pm 0.4	78.1 \pm 0.1	69.2 \pm 0.4	<u>67.2</u> \pm 0.4	90.1 \pm 1.3	<u>81.7</u> \pm 0.5	80.6 \pm 0.9	85.6 \pm 1.1	78.11
Uni-Mol [80]	72.9 \pm 0.6	<u>79.6</u> \pm 0.5	<u>69.6</u> \pm 0.1	65.9 \pm 1.3	91.9 \pm 1.8	82.1 \pm 1.3	<u>80.8</u> \pm 0.3	85.7 \pm 0.2	<u>78.56</u>
KV-PLM [74]	66.9 \pm 1.1	64.7 \pm 1.8	58.6 \pm 0.4	55.3 \pm 0.8	84.3 \pm 1.5	60.2 \pm 2.9	68.8 \pm 4.9	71.9 \pm 2.1	66.29
MoMu [58]	70.5 \pm 2.0	75.6 \pm 0.3	63.4 \pm 0.5	60.5 \pm 0.9	79.9 \pm 4.1	70.5 \pm 1.4	75.9 \pm 0.8	76.7 \pm 2.1	71.63
MoleculeSTM [31]	70.0 \pm 0.5	77.0 \pm 0.4	65.1 \pm 0.3	61.0 \pm 1.0	<u>92.5</u> \pm 1.1	73.4 \pm 2.9	76.9 \pm 1.8	80.8 \pm 1.3	74.57
MolCA [35]	70.0 \pm 0.5	77.2 \pm 0.5	64.5 \pm 0.8	63.0 \pm 1.7	89.5 \pm 0.7	-	-	79.8 \pm 0.5	-
GIT-Mol [30]	73.9 \pm 0.6	75.9 \pm 0.5	66.8 \pm 0.5	63.4 \pm 0.8	88.3 \pm 1.2	-	-	81.1 \pm 1.5	-
BioT5 † [49]	75.2 \pm 0.1	75.5 \pm 0.3	64.4 \pm 0.1	62.4 \pm 0.1	85.8 \pm 0.6	75.9 \pm 0.7	79.7 \pm 0.4	89.2 \pm 0.1	76.01
MV-Mol (w/o view)	73.0 \pm 0.2	79.7 \pm 0.1	69.7 \pm 0.2	65.3 \pm 0.2	93.2 \pm 0.2	81.6 \pm 0.4	79.9 \pm 0.5	87.3 \pm 0.2	78.71
MV-Mol	<u>73.6</u> \pm 0.2	80.3 \pm 0.6	70.0 \pm 0.4	67.3 \pm 0.0	95.6 \pm 1.6	82.1 \pm 0.5	81.4 \pm 0.3	<u>88.2</u> \pm 0.4	79.80

Table 2: Statistics of benchmark datasets. We present the number of molecules, data samples, prediction tasks, and split protocols.

Dataset	# Molecules	# Samples	# Tasks	Split
BBBP	2,039	2,039	1	Scaffold (8/1/1)
Tox21	7,831	7,831	12	
ToxCast	8,597	8,597	617	
SIDER	1,427	1,427	27	
ClinTox	1,478	1,478	2	
MUV	93,807	93,807	17	
HIV	41,127	41,127	1	
BACE	1,513	1,513	1	
PCdes	11,112	11,112	1	Scaffold (7/1/2)
MVST	7,102	16,996	1	
CheBI-20	33,008	33,008	1	Random (8/1/1)

including CheBI [7], PubChem [22], and DrugBank [69], which comprises 273K entities and 643K relations. During pre-training data collection, we exclude molecules within the test set of downstream datasets to alleviate information leakage. More details are presented in Appendix B.

Training Configurations. MV-Mol comprises a molecular structure encoder with 47.3M parameters, a Q-Former with 180.8M parameters, and a text decoder with 140.2M parameters. We pre-train MV-Mol for 70K steps with a batch size of 256 for the modality alignment stage and 50K steps with a batch size of 192 for the knowledge incorporation stage. The pre-training procedure is performed on 4 NVIDIA A100 GPUs for 5 days. We use the AdamW [38] optimizer with a weight decay of 5×10^{-2} . The learning rate is linearly warmed up to 10^{-4} in the first 2K steps and then decreases to 10^{-5} following a cosine annealing strategy. The temperature τ for contrastive learning is fixed as 0.1.

4.2 Molecular Property Prediction (Q1)

To assess the quality of molecular representations, we perform fine-tuning experiments on molecular property prediction, a widely adopted task in MRL.

Datasets. We adopt 8 classification datasets from MoleculeNet [70], a popular benchmark that covers diverse properties of molecules. In line with prior works [32, 80], we adopt Scaffold split with a train/validation/test ratio of 8/1/1. Statistics of the benchmark datasets are presented in Table 2.

Baselines. We adopt two types of baseline MRL models:

- Uni-modal MRL baselines: GraphCL [73], GraphMVP [32], GEM [12] and UniMol [80]. These models are pre-trained solely with molecular structures including 2D topologies, 3D conformations, or both.
- Multi-modal MRL baselines: KV-PLM [74], MoMu [58], MoleculeSTM [31], MolCA [35], GIT-Mol [30], and BioT5 [49]. These models incorporate heterogeneous pre-training data including molecular structures and texts.

Implementation Details. We manually write descriptions for each dataset and feed the prompt and molecular structure into MV-Mol’s encoder to obtain view-based molecular representations. Then, we feed the encoder outputs into a max-pooling layer and a projector for binary classification. The projector is composed of two fully connected layers with ReLU activation. We also implement a variant of our model named MV-Mol (w/o view) by removing view prompts and solely encoding molecules as in Eq. 2. Notably, we apply the prediction head with BioT5’s encoder for a fair comparison of the representation quality. We perform experiments three times with different random seeds and report AUROC scores. Refer to Appendix D for more details of our hyper-parameters and prompts.

Results and Analysis. The overall results of molecular property prediction are displayed in Table 1. The key observations are as follows: (1) Compared with the state-of-the-art model Uni-Mol, MV-Mol achieves an absolute gain of 1.24% on average. Overall, MV-Mol performs best on 5 of 8 datasets and second-best on the remaining 3 datasets. This outstanding performance validates the informativeness of MV-Mol’s molecular representations. (2) On SIDER, ClinTox, and BACE which consist of limited training samples, MV-Mol attains more significant improvements over Uni-Mol (1.4% on SIDER, 3.7% on ClinTox, and 2.5% on BACE), which indicates that adapting molecular representations with view prompts

Table 3: Zero cross-modal retrieval results on the test split of PCdes. The best results are marked in bold, and the second-best results are underlined. \uparrow : the higher the better.

SubTask	Model	MRR \uparrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow
S-T	CLAMP [55]	0.58	0.18	0.49	0.67
	MoMu [58]	11.10	5.93	15.20	20.46
	MoleculeSTM [31]	<u>40.30</u>	<u>26.99</u>	<u>55.33</u>	66.48
	MolCA [35]	33.65	23.21	45.52	54.29
	3D-MoLM [28]	1.68	0.53	1.61	2.57
	MV-Mol (w/o view)	41.17	30.85	53.12	60.23
	MV-Mol	45.25	36.66	55.78	<u>60.41</u>
T-S	CLAMP [55]	0.43	0.08	0.27	0.54
	MoMu [58]	11.58	6.29	15.78	21.86
	MoleculeSTM [31]	30.18	17.45	43.94	<u>57.71</u>
	MolCA [35]	33.48	<u>23.70</u>	<u>44.30</u>	53.30
	3D-MoLM [28]	2.69	0.76	3.60	5.89
	MV-Mol (w/o view)	45.51	35.89	56.68	62.70
	MV-Mol	46.54	37.24	57.17	62.97

Table 4: Zero cross-modal retrieval results on the test split of MVST. The best results are marked in bold, and the second-best results are underlined. \uparrow : the higher the better.

SubTask	Model	MRR \uparrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow
S-T	CLAMP [55]	0.34	0.07	0.23	0.30
	MoMu [58]	13.59	6.97	20.06	26.69
	MoleculeSTM [31]	17.80	9.61	<u>25.98</u>	<u>35.59</u>
	MolCA [35]	<u>18.89</u>	<u>14.25</u>	24.43	27.33
	3D-MoLM [28]	1.09	0.19	1.32	2.22
	MV-Mol (w/o view)	24.87	16.40	34.08	41.32
	MV-Mol	35.38	24.28	47.92	57.27
T-S	CLAMP [55]	0.33	0.04	0.23	0.53
	MoMu [58]	11.81	<u>5.16</u>	17.60	25.26
	MoleculeSTM [31]	<u>12.68</u>	4.86	<u>19.41</u>	<u>29.67</u>
	MolCA [35]	11.05	4.86	18.13	22.43
	3D-MoLM [28]	1.44	0.26	2.00	3.09
	MV-Mol (w/o view)	20.86	9.35	33.37	44.26
	MV-Mol	34.07	22.85	45.73	56.56

brings more significant benefits to low-data scenarios. (3) Removing view descriptions leads to a decline of 0.79% on average. This performance drop corroborates our claims that downstream applications require molecular knowledge from different views, and a universal molecular representation fails to capture the specific context of each task.

4.3 Zero-shot Cross-modal Retrieval (Q2)

We investigate whether MV-Mol captures rich molecular knowledge from heterogeneous sources including texts and knowledge graphs. Hence, we evaluate MV-Mol on cross-modal retrieval, which contains two sub-tasks: structure-to-text (S-T) retrieval and text-to-structure (T-S) retrieval. The former aims to retrieve the most relevant text that describes a given molecule, and the latter aims to retrieve a molecule that best fits the textual description.

Datasets. We incorporate two datasets: PCdes [74] and MVST. The PCdes dataset is collected from the ChEMBL database and comprises 15K molecules with their biochemical definitions and property descriptions. MVST (Multi-View Structure-Text) is a novel dataset introduced by this work. It is collected from chemical, physical, and pharmacokinetic definitions in PubChem. It consists of 7.1K molecules, each corresponding with 2 or more texts from different views. Both datasets are partitioned by Scaffold split with a train/validation/test ratio of 7/1/2. Differing from prior works [58, 74], we perform retrieval on the whole test set and report MRR (mean reversed rank) and Recall at 1/5/10. More details of the datasets are presented in Appendix C.

Baselines. We compare MV-Mol with models that have undergone contrastive learning on molecule-text data including CLAMP [55], MoMu [58], MoleculeSTM [31], MolCA [35] and 3D-MoLM [28].

Implementation Details. On PCdes, we fix the view description as 'biochemical properties and functions'. On MVST, we write descriptions for physical, chemical, and pharmacokinetic views to generate different representations for the same molecule. To further improve the retrieval performance, we modify the re-ranking algorithm in [27] with an ensemble technique. We first retrieve the top- k candidates based on structure-text similarity in Eq. 6. Then, we calculate the CMM logits in Eq. 8 for these k candidates. Finally, we re-rank them by a linear combination of the structure-text similarity scores and the CMM logits.

Results and Analysis. The retrieval results on PCdes are presented in Tab. 3. We observe that: (1) MV-Mol outperforms state-of-the-art methods by 4.95% and 13.06% in MRR for S-T and T-S retrieval. While MoleculeSTM achieves the best results in R@10 on S-T retrieval, it is worth noting that the model is pre-trained on texts similar to those in PCdes and molecules within the test set, raising information leakage concerns. Nevertheless, MV-Mol demonstrates effectiveness in bridging molecular structures and texts. (2) Adding the view description yields a notable gain of 3.08% and 1.03% in MRR for S-T and T-S retrieval, indicating that the view-based molecular representations align better with texts in PCdes.

Table 4 shows the results on MVST, where MV-Mol yields substantial improvements, surpassing MoleculeSTM and MolCA by 16.49% and 21.39% in MRR for S-T and T-S retrieval. The greater performance gain suggests that MV-Mol is endowed with molecular knowledge from a wider range of views. Notably, removing view descriptions leads to a more significant decrease, which we attribute to the multi-view property of the MVST dataset. While a molecule may correspond with multiple descriptions from distinct views, MV-Mol could distinguish between them based on the view-based molecular representations.

4.4 Cross-modal Generation (Q2)

Cross-modal generation also involves two sub-tasks: structure-to-text generation (molecule captioning) and text-to-structure generation (text-based molecule generation).

Datasets. We conduct experiments on ChEMBL-20 [11], a widely-adopted dataset for both tasks. We follow the original split with a train/valid/test ratio of 8/1/1 and adopt evaluation metrics in [10].

Baselines. We compare MV-Mol with the following baselines:

Table 5: Molecule captioning results on the test split of ChEBI-20. 'BL' is short for 'BLEU', and 'R' is short for 'ROUGE'. The best results are marked in bold, and the second-best results are underlined. \uparrow : the higher the better. -: Not reported in the original paper.

Model	#Params	BL-2 \uparrow	BL-4 \uparrow	R-1 \uparrow	R-2 \uparrow	R-L \uparrow	METEOR \uparrow
MolReGPT [26]	-	0.565	0.482	0.623	0.450	0.543	0.585
MolT5-base [10]	250M	0.540	0.457	0.634	0.485	0.578	0.569
MolT5-large [10]	770M	0.594	0.508	0.654	0.510	0.594	0.612
MoMu [58]	770M	0.599	0.515	-	-	0.593	0.597
MolXPT [36]	350M	0.594	0.505	0.660	0.511	0.597	0.626
Text+Chem [6]	250M	0.625	0.542	0.682	0.543	0.622	0.648
MolCA [35]	110M	0.620	0.531	0.681	0.537	0.618	0.651
ChatMol [75]	220M	0.620	0.535	0.677	0.538	0.617	0.644
GIT-Mol [30]	320M	0.352	0.263	0.575	0.485	0.560	0.533
BioT5 [49]	252M	<u>0.635</u>	<u>0.556</u>	<u>0.692</u>	<u>0.559</u>	<u>0.633</u>	<u>0.656</u>
MV-Mol (w/o view)	452M	0.650	0.572	0.698	0.567	0.640	0.669
MV-Mol	452M	0.653	0.575	0.699	0.569	0.640	0.669

- MolReGPT [26]. This baseline leverages a kNN few-shot strategy [44] for in-context learning [9] with ChatGPT [45].
- Molecular language model baselines, including MolT5 [10], MolXPT [36], Text+Chem [6], ChatMol [75] and BioT5 [49] that process 1D SMILES or SELFIES strings of molecules and texts within a unified language model.
- Fusion model baselines, namely MoMu [58]. The baseline connects the GNN outputs with MolT5 to solve molecule captioning.

Implementation Details. For molecule captioning, we fix the view description as 'biochemical properties and functions'. In line with [30, 58], we concatenate the outputs of the MV-Mol encoder with the SELFIES representations of the BioT5 encoder and feed the results into our multi-modal decoder to generate the caption. For text-based molecule generation, we feed the textual representations in Eq. 3 into the MV-Mol decoder to generate the SELFIES string of the molecule.

Results and Analysis. Table 5 shows the results of molecule captioning. Compared to the state-of-the-art model BioT5, MV-Mol shows 1.8% improvements in BLEU scores [47] and 1.3% improvements in METEOR [1] scores, indicating that it generates smoother and more semantically related molecular descriptions.

The performance comparison for text-based molecule generation is reported in Table 5. MV-Mol achieves state-of-the-art performance in most evaluation metrics, highlighted by an outstanding exact ratio of 0.438, surpassing state-of-the-art by 2.5%. While MV-Mol underperforms Text+Chem in fingerprint similarity, we argue that these metrics are only calculated on valid molecules, which may yield over-optimistic results.

Above all, the experiments show that MV-Mol addresses data heterogeneity and can flexibly translate between molecular structures and natural language.

4.5 A Deeper Analysis of View-based Molecular Representations (Q3)

Visualization of View-based Molecular Representations. To investigate how view descriptions affect molecular representations,

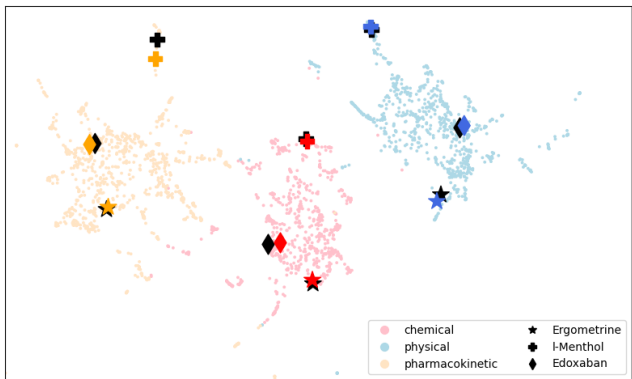


Figure 3: Visualization of view-based molecular representations. We show the molecular representations from chemical, physical, and pharmacokinetic views. We also highlight three molecules and the representations of their textual descriptions from each view.

we visualize the molecular features of MV-Mol on the MVST dataset with UMAP [42]. As shown in Figure 3, the molecular representations based on chemical, physical, and pharmacokinetic views exhibit clear separation along the x-axis. Furthermore, we highlight three representative molecules (indicated by colored shapes) along with corresponding texts from each view (indicated by black shapes), which are closely positioned. Notably, the partial order of the three molecules along the y-axis remains consistent under different views. These observations support our intuition that view-based molecular representations can concurrently capture the consensus and complementary molecular knowledge from different views.

View Prompt Engineering. Motivated by recent advances in prompt engineering [4, 68], we explore how different view prompts affect the representations of our model. We perform analysis on BBBP, SIDER, ClinTox, and BACE from molecule property prediction, and implement the following strategies:

- Empty prompt. This strategy refers to MV-Mol (w/o view), where the view information is not provided.
- Word prompt. We use the name of the dataset as the prompt.
- Sentence prompt. We manually write a brief definition of the prediction task as the prompt.
- Paragraph prompt. We write a well-rounded introduction to the dataset.

As shown in Figure 4b, the sentence prompt works best with BBBP, SIDER, and BACE, while the paragraph prompt yields the best results on ClinTox. We posit that BBBP, SIDER, and BACE focus on specific properties including blood-brain barrier penetration, adverse drug reactions, and target inhibition, where a single sentence suffices to encapsulate all relevant information. In ClinTox, however, more texts are necessary to describe clinical trial criteria and toxicity measures. Besides, the word prompt brings little improvement over the empty prompt. We attribute this to the abbreviated dataset names, which are not frequently used in the pre-training corpora and provide little meaningful information for MV-Mol.

Table 5: Text-based molecule generation results on the test split of ChEBI-20. The best results are marked in bold, and the second-best results are underlined. \uparrow : The higher the better. \downarrow : The lower the better. -: Not reported in the original paper.

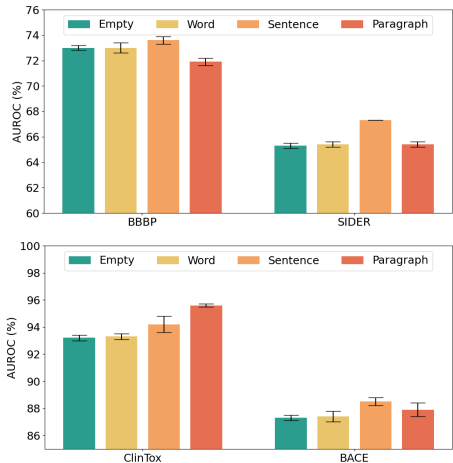
Model	# Params	BLEU \uparrow	Exact \uparrow	Valid \uparrow	Levenshtein \downarrow	MACCS FTS \uparrow	RDKit FTS \uparrow	Morgan FTS \uparrow
MolReGPT [26]	-	0.790	0.139	0.887	24.910	0.847	0.708	0.624
MolT5-base [10]	250M	0.779	0.082	0.786	25.188	0.787	0.661	0.601
MolT5-large [10]	770M	0.854	0.311	0.905	16.071	0.834	0.746	0.684
MoMu [58]	235M	0.815	0.183	0.863	20.520	0.847	0.737	0.678
MolXPT [36]	350M	-	0.215	0.983	-	0.859	0.757	0.667
Text+Chem [6]	250M	0.853	0.322	0.943	16.870	0.901	0.816	0.757
ChatMol [75]	220M	0.848	0.258	0.947	16.759	0.883	0.790	0.726
GIT-Mol [30]	250M	0.756	0.051	0.928	26.315	0.738	0.582	0.519
BioT5 [49]	252M	<u>0.854</u>	<u>0.413</u>	1.000	<u>15.200</u>	0.886	0.801	0.734
MV-Mol	252M	0.858	0.438	1.000	14.952	<u>0.890</u>	<u>0.810</u>	<u>0.745</u>

4.6 Ablation Studies (Q4)

To explore the contribution of our two-stage pre-training, we compare the performances of MV-Mol by removing different pre-training objectives for uni-modal and multi-modal tasks and display the results in Table 6. We observe that: (1) Knowledge graph completion (KGC) loss mainly contributes to cross-modal generation but compromises MV-Mol’s performance on cross-modal retrieval (line 2). (2) Removing the knowledge incorporation stage leads to significant deterioration on all downstream tasks (line 3), highlighting the significance of incorporating multi-view knowledge of molecules. (3) The cross-modal matching objective also benefits all tasks (lines 3 and 4), as it cultivates the joint comprehension of molecules and texts. (4) Removing the modality alignment stage leads to moderate performance drops (line 5), indicating the indispensability of molecular knowledge from large-scale publications. (5) Compared with Stage 1 pre-training, Stage 2 brings more contributions, especially for retrieval tasks, which we attribute to two reasons. First, the knowledge incorporation stage is the key to multi-view MRL, allowing MV-Mol to jointly comprehend view prompts and molecular structures. Second, the knowledge graph enjoys higher quality than the noisy molecule-text pairs for the modality alignment stage.

Table 6: Comparison between MV-Mol with different pre-training objectives. For molecule property prediction, we report the average results on 8 datasets. For zero-shot cross-modal retrieval, we report the average results of Recall at 1/5/10 on the MVST dataset. For cross-modal generation, we report BLEU-2.

Stage 1		Stage 2		Property	Retrieval		Generation S-T
CMC	CMM	KGE	KGC		S-T	T-S	
✓	✓	✓	✓	0.798	0.432	0.417	0.653
✓	✓	✓	✗	0.796	0.437	0.451	0.646
✓	✓	✗	✗	0.786	0.306	0.260	0.647
✓	✗	✗	✗	0.781	0.275	0.248	0.644
✗	✗	✓	✓	0.787	0.395	0.400	0.650

**Figure 4: Experiment results on BBBP, SIDER, ClinTox, and BACE for different prompts.**

5 LIMITATIONS AND BROADER IMPACTS

While our work presents promising results in molecular representation learning, several areas remain for future exploration: (1) Improving the scale and quality of the pre-training data for molecules from structured and unstructured knowledge sources. (2) Exploring the scaling laws of foundation models [20] by incorporating large language models (LLMs) [46, 62] into MV-Mol. (3) Applying MV-Mol to more biomedical entities such as proteins [72], DNA and RNA sequences [54], and cell transcriptomics [78].

MV-Mol bears promise for accelerating biomedical research. However, there is a concern that MV-Mol may be misused to generate potentially dangerous or toxic molecules. Therefore, it is essential to ensure the responsible and ethical use of the model. We emphasize that MV-Mol should be employed solely for research purposes, and any potential medical applications of MV-Mol should undergo comprehensive experimental evaluations.

6 CONCLUSION

In this paper, we present MV-Mol, a molecular representation learning model to harness multi-view molecular expertise from structured and unstructured knowledge sources. To capture the consensus and complementary information across different views, we propose to model views explicitly with textual prompts. We leverage a multi-modal fusion architecture to extract view-based molecular representations, and address the heterogeneity of structured and unstructured knowledge with varying quality and quantity with a two-stage pre-training strategy. Through extensive experiments, we demonstrate the superiority of MV-Mol on molecular property prediction and cross-modal translation. Under thorough analysis aimed at safety, MV-Mol has the potential to deliver unprecedented advancements to the biomedical research community.

ACKNOWLEDGMENTS

This research is supported by the National Key R&D Program of China (No. 2022YFF1203002) and PharMolix Inc.

REFERENCES

- [1] Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*. 65–72.
- [2] Oren M Becker, Alexander D MacKerell Jr, Benoit Roux, and Masakatsu Watanabe. 2001. *Computational biochemistry and biophysics*. Crc Press.
- [3] Guoqing Chao and Shiliang Sun. 2016. Consensus and complementarity based maximum entropy discrimination for multi-view classification. *Information Sciences* 367 (2016), 296–310.
- [4] Banghao Chen, Zhaofeng Zhang, Nicolas Langrené, and Shengxin Zhu. 2023. Unleashing the potential of prompt engineering in Large Language Models: a comprehensive review. *arXiv preprint arXiv:2310.14735* (2023).
- [5] Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. 2020. ChemBERTa: large-scale self-supervised pretraining for molecular property prediction. *arXiv preprint arXiv:2010.09885* (2020).
- [6] Dimitrios Christofidellis, Giorgio Giannone, Jannis Born, Ole Winther, Teodoro Laino, and Matteo Manica. 2023. Unifying molecular and textual representations via multi-task language modelling. *arXiv preprint arXiv:2301.12586* (2023).
- [7] Kirill Degtyarenko, Paula De Matos, Marcus Ennis, Janna Hastings, Martin Zbinden, Alan McNaught, Rafael Alcántara, Michael Darsow, Mickael Guedj, and Michael Ashburner. 2007. ChEBI: a database and ontology for chemical entities of biological interest. *Nucleic acids research* 36, suppl_1 (2007), D344–D350.
- [8] Maxime Delmas, Olivier Filangi, Nils Paulhe, Florence Vinson, Christophe Duperier, William Garrier, Paul-Emeric Saunier, Yoann Pitarch, Fabien Jourdan, Franck Giacomoni, et al. 2021. building a Knowledge Graph from public databases and scientific literature to extract associations between chemicals and diseases. *Bioinformatics* 37, 21 (2021), 3896–3904.
- [9] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234* (2022).
- [10] Carl Edwards, Tuan Lai, Kevin Ros, Garrett Honke, Kyunghyun Cho, and Heng Ji. 2022. Translation between Molecules and Natural Language. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 375–413.
- [11] Carl Edwards, ChengXiang Zhai, and Heng Ji. 2021. Text2mol: Cross-modal molecule retrieval with natural language queries. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 595–607.
- [12] Xiaomin Fang, Lihang Liu, Jieqiong Lei, Donglong He, Shanzhuo Zhang, Jingbo Zhou, Fan Wang, Hua Wu, and Haifeng Wang. 2022. Geometry-enhanced molecular representation learning for property prediction. *Nature Machine Intelligence* 4, 2 (2022), 127–134.
- [13] Yin Fang, Xiaozhuan Liang, Ningyu Zhang, Kangwei Liu, Rui Huang, Zhuo Chen, Xiaohui Fan, and Huajun Chen. 2023. Mol-Instructions-A Large-Scale Biomolecular Instruction Dataset for Large Language Models. In *The Twelfth International Conference on Learning Representations*.
- [14] Yin Fang, Qiang Zhang, Ningyu Zhang, Zhuo Chen, Xiang Zhuang, Xin Shao, Xiaohui Fan, and Huajun Chen. 2023. Knowledge graph-enhanced molecular contrastive learning with functional prompt. *Nature Machine Intelligence* 5, 5 (2023), 542–553.
- [15] Kai Guo, Zhenze Yang, Chi-Hua Yu, and Markus J Buehler. 2021. Artificial intelligence and machine learning in design of mechanical materials. *Materials Horizons* 8, 4 (2021), 1153–1172.
- [16] Linmei Hu, Zeyi Liu, Ziwang Zhao, Lei Hou, Liqiang Nie, and Juanzi Li. 2023. A survey of knowledge enhanced pre-trained language models. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [17] W Hu, B Liu, J Gomes, M Zitnik, P Liang, V Pande, and J Leskovec. 2020. Strategies For Pre-training Graph Neural Networks. In *International Conference on Learning Representations (ICLR)*.
- [18] Ross Irwin, Spyridon Dimitriadis, Jiazhen He, and Esben Jannik Bjerrum. 2022. Chemformer: a pre-trained transformer for computational chemistry. *Machine Learning: Science and Technology* 3, 1 (2022), 015022.
- [19] Pengcheng Jiang, Cao Xiao, Tianfan Fu, and Jimeng Sun. 2023. Bi-level Contrastive Learning for Knowledge Enhanced Molecule Representations. (2023).
- [20] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020).
- [21] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*. 4171–4186.
- [22] Sunghwan Kim, Paul A Thiessen, Evan E Bolton, Jie Chen, Gang Fu, Asta Gindulyte, Lianyi Han, Jane He, Siqian He, Benjamin A Shoemaker, et al. 2016. PubChem substance and compound databases. *Nucleic acids research* 44, D1 (2016), D1202–D1213.
- [23] Nikolaos Kolitsas, Octavian-Eugen Ganea, and Thomas Hofmann. 2018. End-to-End Neural Entity Linking. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*. 519–529.
- [24] Greg Landrum et al. 2013. RDKit: A software suite for cheminformatics, computational chemistry, and predictive modeling. *Greg Landrum* 8 (2013), 31.
- [25] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. (2023).
- [26] Jiatong Li, Yunqing Liu, Wenqi Fan, Xiao-Yong Wei, Hui Liu, Jiliang Tang, and Qing Li. 2023. Empowering Molecule Discovery for Molecule-Caption Translation with Large Language Models: A ChatGPT Perspective. *arXiv preprint arXiv:2306.06615* (2023).
- [27] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems* 34 (2021), 9694–9705.
- [28] Sihang Li, Zhiyuan Liu, Yanchen Luo, Xiang Wang, Xiangnan He, Kenji Kawaguchi, Tat-Seng Chua, and Qi Tian. 2024. Towards 3D Molecule-Text Interpretation in Language Models. *arXiv preprint arXiv:2401.13923* (2024).
- [29] Yingming Li, Ming Yang, and Zhongfei Zhang. 2018. A survey of multi-view representation learning. *IEEE transactions on knowledge and data engineering* 31, 10 (2018), 1863–1883.
- [30] Pengfei Liu, Yiming Ren, and Zhixiang Ren. 2023. Git-mol: A multi-modal large language model for molecular science with graph, image, and text. *arXiv preprint arXiv:2308.06911* (2023).
- [31] Shengchao Liu, Wei Nie, Chengpeng Wang, Jiarui Lu, Zhuoran Qiao, Ling Liu, Jian Tang, Chaowei Xiao, and Animashree Anandkumar. 2023. Multi-modal molecule structure–text model for text-based retrieval and editing. *Nature Machine Intelligence* 5, 12 (2023), 1447–1457.
- [32] Shengchao Liu, Hanchen Wang, Wei Yang Liu, Joan Lasenby, Hongyu Guo, and Jian Tang. 2021. Pre-training Molecular Graph Representation with 3D Geometry. In *International Conference on Learning Representations*.
- [33] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. 2021. Self-supervised learning: Generative or contrastive. *IEEE transactions on knowledge and data engineering* 35, 1 (2021), 857–876.
- [34] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2023. GPT understands, too. *AI Open* (2023).
- [35] Zhiyuan Liu, Sihang Li, Yanchen Luo, Hao Fei, Yixin Cao, Kenji Kawaguchi, Xiang Wang, and Tat-Seng Chua. 2023. MolCA: Molecular Graph-Language Modeling with Cross-Modal Projector and Uni-Modal Adapter. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- [36] Zequn Liu, Wei Zhang, Yingce Xia, Lijun Wu, Shufang Xie, Tao Qin, Ming Zhang, and Tie-Yan Liu. 2023. MolXPT: Wrapping Molecules with Text for Generative Pre-training. *arXiv preprint arXiv:2305.10688* (2023).
- [37] Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel S Weld. 2020. S2ORC: The Semantic Scholar Open Research Corpus. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 4969–4983.
- [38] Ilya Loshchilov and Frank Hutter. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- [39] Yizhen Luo, Kai Yang, Massimo Hong, Xingyi Liu, and Zaiqing Nie. 2023. Molmf: A multimodal molecular foundation model. *arXiv preprint arXiv:2307.09484* (2023).
- [40] Yizhen Luo, Jiahuan Zhang, Siqi Fan, Kai Yang, Yushuai Wu, Mu Qiao, and Zaiqing Nie. 2023. Biomedgpt: Open multimodal generative pre-trained transformer for biomedicine. *arXiv preprint arXiv:2308.09442* (2023).
- [41] Runze Ma, Yidan Zhang, Xinye Wang, Zhenyang Yu, and Lei Duan. 2022. MORN: Molecular Property Prediction Based on Textual-Topological-Spatial Multi-View Learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 1461–1470.
- [42] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. 2018. UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software* 3, 29 (2018), 861.
- [43] David Nadeau and Satoshi Sekine. 2007. A survey of named entity recognition and classification. *Lingvisticae Investigationes* 30, 1 (2007), 3–26.
- [44] Feng Nie, Meixi Chen, Zhirui Zhang, and Xu Cheng. 2022. Improving few-shot performance of language models via nearest neighbor calibration. *arXiv preprint arXiv:2212.02216* (2022).
- [45] OpenAI. 2022. Introducing chatgpt. <https://doi.org/blog/chatgpt>
- [46] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [47] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.
- [48] Qizhi Pei, Lijun Wu, Kaiyuan Gao, Xiaozhuan Liang, Yin Fang, Jinhua Zhu, Shufang Xie, Tao Qin, and Rui Yan. 2024. BioT5+: Towards Generalized Biological Understanding with IUPAC Integration and Multi-task Tuning. *arXiv preprint arXiv:2402.17810* (2024).
- [49] Qizhi Pei, Wei Zhang, Jinhua Zhu, Kehan Wu, Kaiyuan Gao, Lijun Wu, Yingce Xia, and Rui Yan. 2023. BioT5: Enriching Cross-modal Integration in Biology with Chemical Knowledge and Natural Language Associations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 1102–1123.

- [50] Yujia Qin, Yankai Lin, Ryuichi Takanobu, Zhiyuan Liu, Peng Li, Heng Ji, Minlie Huang, Maosong Sun, and Jie Zhou. 2021. ERICA: Improving Entity and Relation Understanding for Pre-trained Language Models via Contrastive Learning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 3350–3363.
- [51] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [52] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. [n.d.]. Improving language understanding by generative pre-training. ([n.d.]).
- [53] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research* 21, 1 (2020), 5485–5551.
- [54] Guillaume Richard, Bernardo P de Almeida, Hugo Dalla-Torre, Christopher Blum, Lorenz Hexemer, Priyanka Pandey, Stefan Laurent, Marie P Lopez, Alexander Laterre, Maren Lang, et al. 2024. ChatNT: A Multimodal Conversational Agent for DNA, RNA and Protein Tasks. *bioRxiv* (2024), 2024–04.
- [55] Philipp Seidl, Andreu Vall, Sepp Hochreiter, and Günter Klambauer. 2023. Enhancing activity prediction models in drug discovery with the ability to understand human language. *arXiv preprint arXiv:2303.03363* (2023).
- [56] Gregory Sliwoski, Sandeepkumar Kothiwale, Jens Meiler, and Edward W Lowe. 2014. Computational methods in drug discovery. *Pharmacological reviews* 66, 1 (2014), 334–395.
- [57] Hannes Stärk, Dominique Beaini, Gabriele Corso, Prudencio Tossou, Christian Dallago, Stephan Günnemann, and Pietro Liò. 2022. 3d infomax improves gnn for molecular property prediction. In *International Conference on Machine Learning*. PMLR, 20479–20502.
- [58] Bing Su, Dazhao Du, Zhao Yang, Yujie Zhou, Jiangmeng Li, Anyi Rao, Hao Sun, Zhiwu Lu, and Ji-Rong Wen. 2022. A molecular multimodal foundation model associating molecule graphs with natural language. *arXiv preprint arXiv:2209.05481* (2022).
- [59] Quan Sun, Qiying Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yuezhe Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. 2023. Generative pretraining in multimodality. *arXiv preprint arXiv:2307.05222* (2023).
- [60] Neil Swainston, Janna Hastings, Adriano Dekker, Venkatesh Muthukrishnan, John May, Christoph Steinbeck, and Pedro Mendes. 2016. libChEBI: an API for accessing the ChEBI database. *Journal of Cheminformatics* 8 (2016), 1–6.
- [61] Paolo Tosco, Nikolaus Stiefl, and Gregory Landrum. 2014. Bringing the MMFF force field to the RDKit: implementation and validation. *Journal of cheminformatics* 6 (2014), 1–4.
- [62] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [63] Murat Ugurlucan, Ilker M Caglar, Fatma N Turhan Caglar, Sedat Ziyade, Oguzhan Karatepe, Yahya Yildiz, Ertugrul Zencirci, Funda Gungor Ugurlucan, Ahmet H Arslan, Semra Korkmaz, et al. 2012. Aspirin: from a historical perspective. *Recent Patents on Cardiovascular Drug Discovery (Discontinued)* 7, 1 (2012), 71–76.
- [64] Yue Wan, Jialu Wu, Tingjun Hou, Chang-Yu Hsieh, and Xiaowei Jia. 2023. From molecules to scaffolds to functional groups: building context-dependent molecular representation via multi-channel learning. *arXiv preprint arXiv:2311.02798* (2023).
- [65] Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021. KEPLER: A unified model for knowledge embedding and pre-trained language representation. *Transactions of the Association for Computational Linguistics* 9 (2021), 176–194.
- [66] Xintao Wang, Qianyu He, Jiaqing Liang, and Yanghua Xiao. 2022. Language models as knowledge embeddings. *arXiv preprint arXiv:2206.12617* (2022).
- [67] Yuyang Wang, Jianren Wang, Zhonglin Cao, and Amir Barati Farimani. 2022. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence* 4, 3 (2022), 279–287.
- [68] Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382* (2023).
- [69] David S Wishart, Yannick D Feunang, An C Guo, Elvis J Lo, Ana Marcu, Jason R Grant, Tanvir Sajed, Daniel Johnson, Carin Li, Zinat Sayeeda, et al. 2018. DrugBank 5.0: a major update to the DrugBank database for 2018. *Nucleic acids research* 46, D1 (2018), D1074–D1082.
- [70] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical science* 9, 2 (2018), 513–530.
- [71] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How Powerful are Graph Neural Networks?. In *International Conference on Learning Representations*.
- [72] Minghao Xu, Xinyu Yuan, Santiago Miret, and Jian Tang. 2023. Protst: Multi-modality learning of protein sequences and biomedical texts. In *International Conference on Machine Learning*. PMLR, 38749–38767.
- [73] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph contrastive learning with augmentations. *Advances in neural information processing systems* 33 (2020), 5812–5823.
- [74] Zheni Zeng, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2022. A deep-learning system bridging molecule structure and biomedical text with comprehension comparable to human professionals. *Nature communications* 13, 1 (2022), 862.
- [75] Zheni Zeng, Bangchen Yin, Shipeng Wang, Jiarui Liu, Cheng Yang, Haishen Yao, Xingzhi Sun, Maosong Sun, Guotong Xie, and Zhiyuan Liu. 2023. Interactive Molecular Discovery with Natural Language. *arXiv preprint arXiv:2306.11976* (2023).
- [76] Ningyu Zhang, Zhen Bi, Xiaozhuan Liang, Siyuan Cheng, Haosen Hong, Shumin Deng, Qiang Zhang, Jiazhang Lian, and Huajun Chen. 2021. OntoProtein: Protein Pretraining With Gene Ontology Embedding. In *International Conference on Learning Representations*.
- [77] Haiteng Zhao, Shengchao Liu, Ma Chang, Hannan Xu, Jie Fu, Zhihong Deng, Lingpeng Kong, and Qi Liu. 2024. Gimlet: A unified graph-text model for instruction-based molecule zero-shot learning. *Advances in Neural Information Processing Systems* 36 (2024).
- [78] Suyuan Zhao, Jiahuan Zhang, Yizhen Luo, Yushuai Wu, and Zaiqing Nie. 2024. LangCell: Language-Cell Pre-training for Cell Identity Understanding. *arXiv preprint arXiv:2405.06708* (2024).
- [79] Qinghai Zheng, Jihua Zhu, Zhongyu Li, Zhiqiang Tian, and Chen Li. 2023. Comprehensive multi-view representation learning. *Information Fusion* 89 (2023), 198–209.
- [80] Gengmo Zhou, Zhifeng Gao, Qiankun Ding, Hang Zheng, Hongteng Xu, Zhewei Wei, Linfeng Zhang, and Guolin Ke. 2023. Uni-Mol: a universal 3D molecular representation learning framework. (2023).
- [81] Jinhua Zhu, Yingce Xia, Lijun Wu, Shufang Xie, Tao Qin, Wengang Zhou, Houqiang Li, and Tie-Yan Liu. 2022. Unified 2d and 3d pre-training of molecular representations. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2626–2636.
- [82] Jinhua Zhu, Yingce Xia, Lijun Wu, Shufang Xie, Wengang Zhou, Tao Qin, Houqiang Li, and Tie-Yan Liu. 2023. Dual-view molecular pre-training. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3615–3627.

A DETAILS OF MV-MOL

The representation similarity $\text{sim}_{\text{tri}}(\cdot, \cdot)$ in Eq. 11 is calculated as follows:

$$\begin{aligned} \text{sim}_{\text{tri}}(z^{(h,r)}, z^{(t)}) &= \begin{cases} \max_{i=1,2,\dots,K} \left\{ f_{\text{proj}}(z_i^{(h,r)})^T f_{\text{proj}}(z_1^{(t)}) \right\}, & h \in E_{\text{mol}}, t \in E_{\text{text}}, \\ \max_{i,j=1,2,\dots,K} \left\{ f_{\text{proj}}(z_i^{(h,r)})^T f_{\text{proj}}(z_j^{(t)}) \right\}, & h, t \in E_{\text{mol}}, \\ f_{\text{proj}}(z_1^{(h,r)})^T f_{\text{proj}}(z_1^{(t)}), & h, t \in E_{\text{text}} \end{cases} \end{aligned} \quad (\text{A1})$$

B PRE-TRAINING DATA COLLECTION

Following [58, 74], we perform named entity recognition (NER) [43] and entity linking (EL) [23] to obtain molecule-text pairs from biomedical literature. Specifically, we first sample 321K frequently accessed molecules from PubChem [22]. Then we identify mentions of these molecules within 3.5M publications from S2ORC [37]. We start with the corresponding sentence of the mention and randomly extend the context by the previous or next sentence within the paragraph until the text snippet exceeds 256 tokens. Consequently, we collect 12M molecule-text pairs for 60K molecules.

We also collect a knowledge graph for molecules by combining public databases including PubChem [22], ChEBI [7], DrugBank [69] and FORUM [8]. We first extract 196K molecules from these databases and perform deduplication. Then, we collect the connections between these molecules and other entities as follows:

Table A1: Statistics of the collected knowledge graph. The left column denotes the entity or relation type, and the right column denotes the number of entities and relations.

Entities	
Molecules	196,454
Targets	403
Ontologies	22,076
Diseases	2,835
Properties	51,882
All	273,650
Relations	
Molecule-Target Interaction	23,870
Molecule-Ontology Connection	252,555
Molecule-Disease Association	172,241
Molecule-Physical Property	6,166
Molecule-Chemical Property	38,206
Molecule-Pharmacokinetic Property	7,510
Molecule-Molecule Connection	119,793
Ontology-Ontology Connection	22,706
All	643,047

- We collect drug property descriptions that span diverse views from PubChem. We treat each description as a unique node and connect it with the corresponding molecule.
- We collect drug targets, drug enzymes, drug carriers, and drug transporters from DrugBank. We also collect BindingDB compounds and protein targets with $K_i \leq 10nM$.
- We collect drug-disease and drug-ontology connections from FORUM with $q_value < 10^{-6}$.
- We collect drug-ontology and ontology-ontology connections from CheBI with libCheBI toolkit [60].

In total, the knowledge graph comprises 273K entities and 643K relationships. Table A1 presents the overall statistics of the knowledge graph. Following [80], we calculate 3D conformations by the MMFF algorithm [61] within the RDKit [24] package for molecules.

C BENCHMARK DATASETS

C.1 Molecular Property Prediction

We adopt 8 classification datasets for molecular property prediction, including BBBP (Blood-Brain Barrier Penetration), Tox21 (Toxicology in the 21st Century), ToxCast (Toxicity), SIDER (Side Effects Resource), ClinTox (Clinical Trial Outcomes and Toxicity), MUV (Maximum Unbiased Validation), HIV (HIV inhibition), and BACE (Human β -secretase 1 Binding).

C.2 Cross-Modal Retrieval

For cross-modal retrieval, we incorporate the following datasets: **PCDes**. The original dataset [74] comprises 15K molecules and their descriptions in CheBI. We filter out 3,888 molecules that have appeared in our pre-training dataset. Unlike the original experiment

setting, we adopt the Scaffold split with a train/validation/test ratio of 7/1/2 and perform zero-shot retrieval on the test set with the entire paragraph.

MVST. This dataset is introduced as part of our work. Different from PCDes, each molecule within MVST corresponds to multiple descriptive texts, distinguished by a given view. The dataset primarily focuses on 3 types of views:

- The chemical view. The data for this view is directly derived from the CheBI description of molecules within the PubChem database.
- The physical view. We assemble multiple columns within PubChem including the general physical description, color, form, order, boiling point, melting point, flashing point, solubility, density, vapor pressure, stability (shelf life), and decomposition. The text is organized as '*COLUMN_NAME: COLUMN_DESCRIPTION*'. We randomly sample rows if the assembled text exceeds 256 tokens.
- The pharmacokinetic view. We also incorporate multiple columns from PubChem including drug indication, disposal methods, drug-food interactions, toxicity summary, metabolism, therapeutic uses, inhalation systems, absorption, distribution, exertion, biological half-life, adverse effects, and carcinogenicity. The organization of the text and length control follows the same as the physical view.

Following previous works [10, 74], we replace the name of the molecule within texts with '*the molecule*'. After collecting data from each view, we sample molecules with textual descriptions from 2 or more views and partition them by Scaffold split with a train/validation/test ratio of 7/1/2.

C.3 Cross-Modal Generation

We incorporate the CheBI-20 dataset, a widely adopted dataset collected from the CheBI database. We follow the original random split with a train/validation/test ratio of 8/1/1.

D DOWNSTREAM EXPERIMENT DETAILS

We manually write descriptions for each dataset in MoleculeNet as view prompts to improve molecular representations. Specifically:

- BBBP. The prompt is '*blood-brain barrier penetration (permeability)*'.
- Tox21. The prompt is '*Qualitative toxicity measurements including nuclear receptors and stress response pathways*'.
- ToxCast. The prompt is '*Qualitative toxicity measurements*'.
- SIDER. The prompt is '*adverse drug reactions (ADR) for 27 system organ classes*'.
- ClinTox. The prompt is '*Qualitative data of drugs if they failed clinical trials for toxicity reasons*'.
- MUV. The prompt is '*Subset of PubChem BioAssay designed for validation of virtual screening techniques*'.
- HIV. The prompt is '*Experimentally measured abilities to inhibit HIV replication*'.
- BACE. The prompt is '*Binding results for human β -secretase 1 (BACE-1)*'.