

From Artificially Real to Real: Leveraging Pseudo Data from Large Language Models for Low-Resource Molecule Discovery

Yuhan Chen, Nuwa Xi, Yanrui Du, Haochun Wang, Jianyu Chen, Sendong Zhao*, Bing Qin

Research Center for Social Computing and Information Retrieval
Harbin Institute of Technology, China
{yuhanchen, sdzhao}@ir.hit.edu.cn

Abstract

Molecule discovery serves as a cornerstone in numerous scientific domains, fueling the development of new materials and innovative drug designs. Recent developments of in-silico molecule discovery have highlighted the promising results of cross-modal techniques, which bridge molecular structures with their descriptive annotations. However, these cross-modal methods frequently encounter the issue of data scarcity, hampering their performance and application. In this paper, we address the low-resource challenge by utilizing artificially-real data generated by Large Language Models (LLMs). We first introduce a retrieval-based prompting strategy to construct high-quality pseudo data, then explore the optimal method to effectively leverage this pseudo data. Experiments show that using pseudo data for domain adaptation outperforms all existing methods, while also requiring a smaller model scale, reduced data size and lower training cost, highlighting its efficiency. Furthermore, our method shows a sustained improvement as the volume of pseudo data increases, revealing the great potential of pseudo data in advancing low-resource cross-modal molecule discovery.

Introduction

Molecule discovery plays a critical role in numerous scientific domains including chemistry (Wang et al. 2023b; Cuzzucoli Crucitti et al. 2023), pharmacology (Patani and LaVoie 1996; Anderson 2003), and materials science (Curtarolo et al. 2013). However, traditional molecule design methods are frequently faced with challenges such as high costs, lengthy development processes, and limited success rates. Introducing a new drug to the market, for instance, might demand over a billion dollars and more than a decade of development (Gaudelet et al. 2021).

With the advent of artificial intelligence (AI), innovative cross-modal methods are ushering in new ways to synthesize and analyze complex molecular structures, enhancing efficiency and reshaping the fields of computational chemistry and material science. Edwards et al. (2022) proposed a novel approach to directly translate molecules to corresponding captions and generate molecular structures from natural language text, shown in Figure 1. This cross-modal

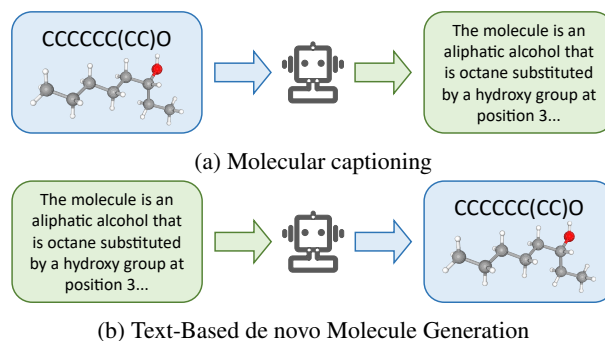


Figure 1: Illustration of translation between molecule and description in cross-modal molecule discovery.

method heralds a future in which the design and study of specialized molecules can be achieved through simple natural language sentences.

Various attempts have been made to resolve these tasks. Molt5 (Edwards et al. 2022) uses SMILES (Simplified Molecular Input Line Entry System) (Weininger 1988) and molecule description respectively for masked language modeling (MLM) (Raffel et al. 2020) as pre-training. Liu et al. (2023) pre-train models with causal language modeling (CLM) on the sequences that blend biomedical literature with molecular structural representations, derived from replacing molecular entities with their SMILES representations. However, these studies are limited by the scarcity of parallel molecule-description pairs, rendering direct sequence-to-sequence training unfeasible. The effectiveness of sequence-to-sequence (seq2seq) training is evident in Christofidellis et al. (2023), where the annotated data from the downstream dataset is incorporated for pre-training, albeit in a significantly lower ratio compared to the unannotated data. The primary bottleneck is the annotation process itself: the annotation of these pairs demands specialized knowledge in molecular chemistry, rendering large-scale human annotation both expensive and difficult.

Inspired by the great success of LLMs in natural language processing (NLP) and related fields (Bagal et al. 2021; Frey et al. 2022; Ferruz, Schmidt, and Höcker 2022), we propose to mitigate the low-resource difficulty by using artificially-real data generated by LLMs. Unlike “real data”, which orig-

*Corresponding author

inates from genuine experimental or observational sources, this “pseudo data” or “artificially-real data” is crafted artificially. While it mirrors the format of real data, its content does not depict actual real-world observations, making it potentially unsuitable for direct real-world applications.

Our approach begins by creating a comprehensive pseudo dataset intended for seq2seq pre-training. We collect 1M unlabeled molecules from PubChem and use the in-context learning ability of LLMs to generate descriptive captions for these molecules. To ensure the integrity and diversity of this pseudo data, we adopt a retrieval-based one-shot prompting strategy during generation. Through this way, we construct the first artificially-real dataset, PseudoMD-1M, consisting of 1,020,139 pseudo molecule-description pairs.

Based on this dataset, we explore the optimal method to leverage pseudo data. We propose two primary methods: 1) using pseudo data exclusively during pre-training for domain adaptation, and 2) integrating pseudo data with real data during fine-tuning as a data augmentation technique. To offer a comprehensive evaluation, we further compile DrugBank-23, a novel dataset derived from a different data source than existing datasets.

In summary, our contributions are as follows:

- We are the first to incorporate LLMs for low-resource molecule discovery. Using artificially-real data generated by LLMs, we are able to mitigate the data scarcity for the tasks. We release PseudoMD-1M, the first artificially-real dataset for cross-modal molecule discovery, which is $33\times$ larger than existing real datasets.
- We explore the effective construction and utilization of pseudo data. We specifically investigate two principal techniques, including using pseudo data as domain adaptation and data augmentation. We conduct comprehensive experiments on existing datasets, and provide our new dataset called DrugBank-23, which adds a novel data source compared to current datasets.
- Experimental results show that despite smaller model size and amount of pre-training data, models using artificially-real data as domain adaptation outperform all prior methods. Furthermore, our method shows continuous improvement with increasing volumes of pseudo data, underscoring its promising future applications.

Related Work

Cross-Modal Molecule Discovery

With the advancement of in-silico molecule discovery methods, the field of molecule exploration is undergoing a transformative shift away from its resource-intensive and costly origins (Rifaioglu et al. 2019; Gaudelet et al. 2021). Edwards, Zhai, and Ji (2021) introduce a new task Text2Mol, which uses descriptions as search queries to retrieve the target molecules. Following this, Edwards et al. (2022) propose two innovative tasks: molecule captioning and text-guided de novo molecule generation. These tasks aim at translating between molecular structures and natural language texts. MolXPT (Liu et al. 2023) leverages literature annotations of molecules to construct a pre-training dataset. Christofidelis et al. (2023) further improves the field with multi-task

learning, which combines single-domain and cross-domain datasets for joint training. Most recently, Li et al. (2023) propose a strategy that enables LLMs to accomplish both molecule captioning and text-guided molecule generation tasks. Here we take one step further to construct a large number of high-quality parallel data pairs, in response to the data scarcity that limits the performance of the above approaches.

Large Language Models

LLMs have achieved significant success in natural language processing by scaling up to billions of parameters (Brown et al. 2020; Ouyang et al. 2022). Trained on vast corpora (Singhal et al. 2023), LLMs show more general intelligence (Bubeck et al. 2023) and remarkable capabilities such as in-context learning (Rubin, Herzig, and Berant 2022; Min et al. 2022). They have also obtained promising performance in chemical (Bagal et al. 2021; Frey et al. 2022), biological (Ferruz, Schmidt, and Höcker 2022; Xi et al. 2023) and medical (Wang et al. 2023a; Du et al. 2023) domains. Due to their great generation capability, numerous works have relied on LLMs to generate data for various purposes, including creating semantic textual similarity datasets (Schick and Schütze 2021), augmenting natural language inference (Liu et al. 2022), automatically formulating instructions (Wang et al. 2022) and improving few-shot retrieval (Dai et al. 2022). Inspired by these achievements, we aim to employ LLMs to generate parallel data, addressing data scarcity in cross-modal molecule discovery.

Methodology

Task Overview

Here we introduce two primary tasks for cross-modal molecule discovery. First proposed by Edwards et al. (2022), the two tasks act as a bridge between molecule discovery and NLP and can be considered as cross-modal translation tasks.

Molecular Captioning As illustrated in Figure 1a, Given the SMILES representation $S_{\mathcal{M}}$ of molecule \mathcal{M} , the task is to generate the corresponding descriptions $\mathcal{D}_{\mathcal{M}}$.

Text-Based De Novo Molecule Generation As shown in Figure 1b, given the descriptions $\mathcal{D}_{\mathcal{M}}$ of molecules \mathcal{M} , the task is to generate its corresponding SMILES $S_{\mathcal{M}}$.

Artificially-Real Data Generation

High-quality pseudo data is the foundation for further exploration. Here we propose PseudoMD-1M, the first pseudo dataset composed of 1M parallel molecule-description data pairs. To acquire sufficient data, we leverage a vast number of unlabeled molecules and use LLMs to generate corresponding descriptions. We begin by collecting 1.1 million unannotated SMILES strings of molecules from PubChem (Kim et al. 2023). We then employ a rigorous filtering procedure to filter out the SMILES in downstream datasets to ensure that there is no overlap between the collected molecules and those contained in the real datasets (Edwards, Zhai, and Ji 2021; Zeng et al. 2022). By doing so, we ensure that no supplementary information about the molecules present in

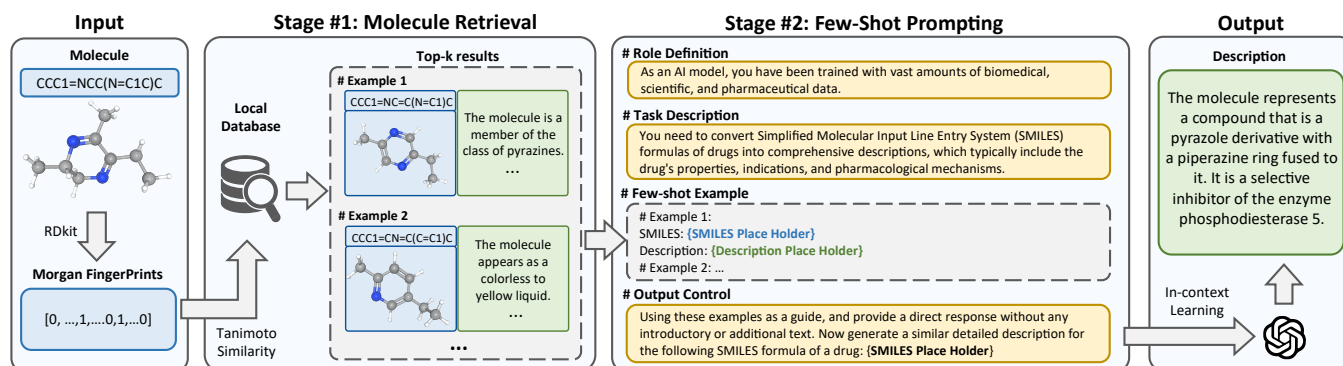


Figure 2: The workflow for pseudo data generation. Starting with an unlabeled molecule represented by its Morgan Fingerprints, two stages are involved. In stage 1, the input molecule serves as a search query to retrieve the top-k similar molecules from a local database containing 37,898 annotated molecule-caption pairs. In stage 2, the retrieved molecules and their captions are integrated into a prompt. Then LLMs perform in-context learning and generate a description for the input molecule.




Figure 3: Comparison of data quality. We use the method proposed by Edwards et al. (2022) to evaluate the similarity between molecule-description pairs as an estimation of the data quality. The distribution is visualized using Kernel Distribution Estimation. A higher Text2Mol score signifies closer molecule-description resemblance, and “Density” represents the data concentration in a given region.

the real datasets is accidentally incorporated, thereby maintaining the integrity and independence of the training process. With ChatGPT API, we generate textual descriptions that encompass key aspects such as properties and structural features for each unannotated molecule. To improve the quality of generated descriptions, we implement a retrieval-based prompt paradigm that comprises two main stages as follows: Molecule Retrieval and Few-Shot Prompting.

Molecule Retrieval In-context learning (Brown et al. 2020) is one of the emergent abilities of LLMs, and the instances used in the prompts given to the LLMs play an important role in the generation quality. As molecules with similar structures often display corresponding characteristics (Wang et al. 2016), we retrieve the descriptions of annotated molecules that resemble the unlabeled molecule, using them as the few-shot instance during prompting. Specifically, we collect 37,898 annotated molecules with captions

from PubChem (Kim et al. 2023), then retrieve the molecules with top-k Tanimoto similarity (Tanimoto 1958), a standard measure in cheminformatics. To prevent information leakage during testing, we exclude the molecules that are contained in the real data test set (Edwards, Zhai, and Ji 2021; Zeng et al. 2022). This process enables the models to learn from the information embedded within the descriptions of molecules that possess similar properties, ensuring a more tailored and accurate representation. Figure 3 shows the estimate of the data quality, indicating that the few-shot prompting approach (in blue) yields higher-quality data, more closely resembling real data than without.

Few-Shot Prompting Upon retrieving the top-k results for each unlabeled molecule from our local database, we select one example using a weighted distribution, where molecules with higher similarity have a greater chance of being chosen. This selected example is then incorporated into the final prompt. We opt for one-shot prompting to minimize generation costs, as expenses increase linearly with the number of instances included in few-shot prompts. This weighted selection method prevents repetitive selection of the same molecule as the few-shot example, thereby improving the diversity during generation while maintaining the similarity between the molecule to be annotated and the few-shot example. As shown in Figure 2, the complete prompt comprises role definition, task description, few-shot example, and output control. The role definition and task description give LLMs the general context and enable its learned knowledge, while the few-shot example acts like a supplementary material for the LLMs to refer to. Then, with the output control for format clarification, the LLMs should be able to generate the desired description.

Approaches to Utilize Artificially Real Data

The ways to utilize the pseudo data decide how the model will perform on real data. We propose and explore two primary strategies to optimize the use of pseudo data.

Pseudo Data as Data Augmentation Data augmentation strategy can be roughly categorized into two kinds, modifi-

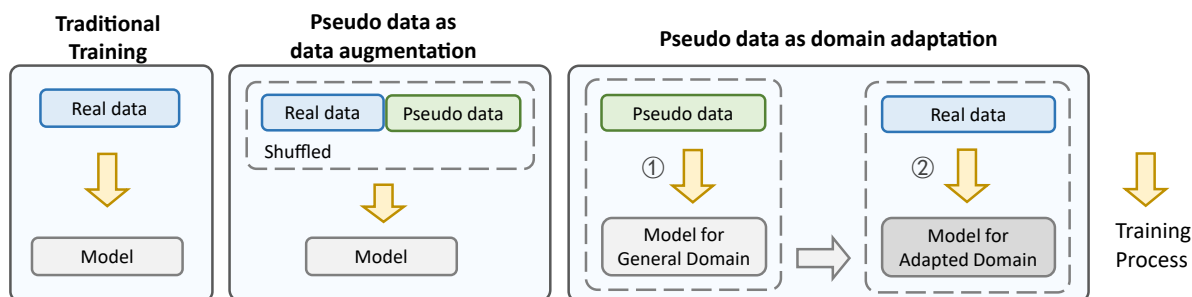


Figure 4: Different methods for utilizing pseudo data. Traditional training employs only the real dataset for fine-tuning. The data augmentation approach fine-tunes the model on the combined dataset with pseudo data incorporated. In the domain adaptation method, the model is (1) initially pre-trained on two concurrent cross-modal translation tasks using pseudo data as domain adaptation, and (2) further trained on each task using real data.

cation of existing data and generation of pseudo data. The former takes an existing data instance and makes certain alterations to it without changing its inherent meaning or label, such as rotation, flipping, and cropping for images (Krizhevsky, Sutskever, and Hinton 2012), or synonym replacement for text (Wang and Yang 2015; Wei and Zou 2019; Miao et al. 2020). This method is more about adding variability and noise to existing data instances than generating completely new ones. The latter, on the other hand, involves creating new data instances that did not exist in the original dataset based on the characteristics and distribution of the original data, which is an efficient alternative when real data is scarce or when creating new real data is costly or unfeasible. Existing applications include back translation for text (Sennrich, Haddow, and Birch 2016), and GANs for images (Goodfellow et al. 2014).

Inspired by the latter techniques, we explore the use of pseudo data as data augmentation. As shown in Figure 4, we keep the original data in the training set and augment them with pseudo data during fine-tuning. Using the same method as described in Figure 3, we assess the distribution of the real training set and the sample the augmented pseudo data based on the same distribution, ensuring consistency in the overall dataset distribution before and after data augmentation. We hope that this data augmentation approach using pseudo data will expose the model to a broader range of data patterns and scenarios, thus enhancing its ability to recognize complex patterns and generalize its learning to unseen data.

Pseudo Data as Domain Adaptation Models pre-trained on general domain might perform less ideally when it is applied to specific domains for which they were not explicitly trained (Malte and Ratadiya 2019). In our case, the SMILES appears as an unfamiliar symbol to such models, making the direct fine-tuning approach less efficient. To bridge this gap, we use pseudo data as a second pre-training stage for domain adaptation. As shown in Figure 4, we train the model using pseudo data for two concurrent cross-modal translation tasks: molecular captioning and text-based de novo molecule generation. Using a direct and bidirectional seq2seq approach, this stage is intended to empower the model to not only recognize the SMILES representation but also to grasp the relationship between natural

language and SMILES. Given that our primary focus at this stage is not on data authenticity, pseudo data emerges as a preferable choice, particularly because it provides a large number of parallel data pairs for supervised seq2seq training compared to real datasets. We then further fine-tune it on real data to refine and enhance the model’s understanding of SMILES for further authenticity – a critical aspect for applications like drug discovery.

Experiments

To validate the effectiveness of using pseudo data, we conduct comprehensive experiments comparing our proposed approaches with existing methods. We further conduct experiments to demonstrate how the balance between real data and pseudo data could affect model performance. All the experiments are conducted on both molecular captioning and molecule generation. The implementation details are listed in Appendix C.

Settings

Datasets Currently, only a few datasets with parallel molecule-description pairs exist, including ChEBI-20 (Edwards, Zhai, and Ji 2021) and PCdes (Zeng et al. 2022), both constructed using data from PubChem (Kim et al. 2023). To enhance evaluation comprehensiveness, we assemble a new dataset called DrugBank-23, based on DrugBank (Wishart et al. 2018). We experiment on all three datasets (ChEBI-20, PCdes, and DrugBank-23). The detailed information about these datasets is listed in Table 1.

Models We evaluate the following methods:

- **T5** (Raffel et al. 2020). T5 directly fine-tuned on downstream datasets.
- **MolT5** (Edwards et al. 2022). T5 pre-trained with MLM using SMILES and molecule descriptions respectively, then fine-tuned on downstream datasets.
- **ChatGPT** (Li et al. 2023). GPT-3.5-Turbo using few-shot prompting strategy. We cite the results from the original paper on ChEBI-20, then apply the same strategy to test on the other datasets.

Info	ChEBI-20	PCdes	DrugBank-23
Train	26,407	10,500	17,109
Validation	3,301	1,500	3,667
Test	3,300	3,000	3,666
$\mathcal{L}_{\text{SMILES}}$	81.56	56.47	54.11
$\mathcal{L}_{\text{Description}}$	52.88	72.47	65.04
Data source	PubChem		DrugBank

Table 1: Details about the existing datasets and ours (DrugBank-23). $\mathcal{L}_{\text{SMILES}}$ denotes the average length of SMILES while $\mathcal{L}_{\text{Description}}$ denotes the average word count per description.

Model	Data scale	Steps	Backbone
MolT5	500M	1M	T5 _{large}
MolXPT	8M	200k	GPT2 _{medium}
Text&Chem T5	33.5M	131k	T5 _{base}
Aug-T5	0	0	T5 _{small}
Ada-T5	1M	100k	T5 _{small}

Table 2: Pre-training details for different Models. ‘‘M’’ stands for million and ‘‘k’’ denotes thousand.

- **MolXPT** (Liu et al. 2023). GPT-2 pre-trained with CLM using abstracts of biomedical literature where molecules are replaced with the corresponding SMILES, then fine-tuned on downstream datasets. As the model is currently unavailable, we cite their results on ChEBI-20.
- **Text&Chem T5** (Christofidellis et al. 2023). T5 pre-trained using multi-task learning, then fine-tuned on downstream datasets.
- **Aug-T5** (ours). T5 fine-tuned on datasets augmented with pseudo data from PseudoMD-1M, sampled from 1k to 512k, doubling at each step. We report the optimal performances for each dataset. See Appendix D for details.
- **Ada-T5** (ours). T5 pre-trained using molecule-description pairs from PseudoMD-1M as domain adaptation, then fine-tuned on downstream datasets.

As shown in Table 2, both our proposed methods utilize the smallest model scale, pre-training data, and steps, while Aug-T5 requires no additional pre-training. We first test our methods on T5_{small} (Aug-T5/Ada-T5) and then apply them to T5_{base} (Aug-T5_{base}/Ada-T5_{base}).

Metrics Following existing studies (Edwards et al. 2022; Liu et al. 2023; Christofidellis et al. 2023), we evaluate the results for molecular captioning with BLEU-2, BLEU-4 (Papineni et al. 2002), ROUGE-1, ROUGE-2, ROUGE-L (Lin 2004) and METEOR (Banerjee and Lavie 2005), and BLEU-4 (Papineni et al. 2002), Accuracy (Edwards et al. 2022), Validity (Polykovskiy et al. 2020), Levenshtein distance (Miller, Vandome, and McBrewster 2009), MACCS-FTS (Durant et al. 2002), RDKit-FTS (Schneider, Sayle, and Landrum 2015), Morgan-FTS (Rogers and Hahn 2010) and FCD (Preuer et al. 2018) for text-based de novo molecule

generation. Selected metrics are presented in Tables 3, 4 and Figures 5 and 6, with comprehensive results in Appendix D.

Comparison with Existing Methods

Results on Molecular Captioning Table 3 shows the results of different models for molecule captioning. Ada-T5 outperforms all previous methods and achieves the state-of-the-art on all three datasets across all the metrics. Compared to the previous state-of-the-art, Ada-T5 uses less than 3% of the pre-training data and only a third of the model parameters, yet requires fewer training steps, demonstrating the effectiveness and computational efficiency of high-quality pseudo data. On the other hand, Aug-T5 outperforms T5, MolT5, ChatGPT and has comparable performance with MolXPT and Text&Chem T5, using 9%-30% of the parameters and requires no pre-training. This highlights the benefit from the enhanced diversity of descriptions by incorporating pseudo data into the training set. Meanwhile, Ada-T5_{base} makes an extra but relatively little progress compared to Ada-T5, indicating that although using pseudo data for domain adaptation could also benefit from the expansion of model size like most methods, the exploitation of pseudo data only demands a relatively small number of parameters. In contrast, Aug-T5_{base} mirrors the results of its smaller version, indicating that for data augmentation, simply increasing the model scale may not offer substantial benefits. One thing to notice is that despite the data used to train the model is generated by ChatGPT API, both our trained models can still beat ChatGPT across different metrics. This indicates that although ChatGPT can accomplish the task to a certain extent, the data it generated can still help the models achieve a more seamless transition through pre-training from general domain to this domain.

Results on Text-Based Molecule Generation Table 4 presents the results of different models for molecule generation. Ada-T5 achieves the best performance in all three datasets across almost all metrics, demonstrating its capability to generate high-quality SMILES. The only exception is that the MolXPT slightly surpasses Ada-T5 by 0.009 in ChEBI-20 dataset on the validity metric, which is calculated using RDKit to simply check whether the string can be successfully converted to a molecule object without errors and whether the molecule represents a realistic and feasible chemical structure, without any comparison to the targeted SMILES and the input descriptions. Despite this one slight superiority, MolXPT performs significantly worse than Ada-T5 on other metrics, meaning that although it can generate slightly more valid SMILES, it does not take into account the designated instructions, ergo making it one step away from real-world application.

On the other hand, Aug-T5 surpasses some existing methods in certain datasets on specific metrics. However, its consistency falls short compared to Ada-T5. This variability may be traced back to the construction of molecule-description data pairs in pseudo data: the LLMs use the real SMILES as the input, leaving only the description part of the pseudo data genuinely ‘‘pseudo’’. This means that when training Aug-T5 on molecule captioning, it gets the

Model	Parameters	ChEBI-20			PCdes			DrugBank-23		
		BL	RG	MET	BL	RG	MET	BL	RG	MET
T5	800M	0.467 ^{†*}	0.478 ^{†*}	0.586 ^{†*}	0.252 ^{†*}	0.259 ^{†*}	0.367 ^{†*}	0.272 ^{†*}	0.299 ^{†*}	0.396 ^{†*}
MolT5	800M	0.508 [†]	0.510 ^{†*}	0.614 [†]	0.266 [†]	0.272 [†]	0.380 ^{†*}	0.293 [†]	0.317 [†]	0.416 [†]
MolXPT	350M	0.505 ^{†*}	0.511 ^{†*}	0.626 [†]	-	-	-	-	-	-
Text&Chem T5	250M	0.542 [†]	0.543 [†]	0.648 [†]	0.266 [†]	0.274 [†]	0.382 [†]	0.280 ^{†*}	0.312 ^{†*}	0.413 ^{†*}
ChatGPT	-	0.482 ^{†*}	0.450 ^{†*}	0.585 ^{†*}	0.194 ^{†*}	0.193 ^{†*}	0.315 ^{†*}	0.191 ^{†*}	0.218 ^{†*}	0.325 ^{†*}
Aug-T5	77M	0.515	0.517	0.621	0.270	0.275	0.385	0.297	0.322	0.421
Aug-T5 _{base}	250M	0.516	0.520	0.620	0.268	0.272	0.383	0.294	0.316	0.416
Ada-T5	77M	0.553	0.552	0.652	0.295	0.295	0.406	0.310	0.337	0.435
Ada-T5 _{base}	250M	0.564	0.562	0.660	0.295	0.297	0.409	0.322	0.346	0.445

Table 3: Results of different models for molecular captioning on ChEBI-20, PCdes and DrugBank-23 datasets. [†]/^{*} denotes that Ada-T5_{base}/Aug-T5_{base} perform significantly better than baselines at p -value < 0.01 using t-test. The best scores are in bold. BL: BLEU-4. RG: ROUGE-2. MET: METEOR.

Model	Parameters	ChEBI-20			PCdes			DrugBank-23		
		Acc	Val	MAC	Acc	Val	MAC	Acc	Val	MAC
T5	800M	0.279 ^{†*}	0.902 ^{†*}	0.823 ^{†*}	0.089 [†]	0.910 ^{†*}	0.698 [†]	0.131 ^{†*}	0.923 ^{†*}	0.682 [†]
MolT5	800M	0.311 ^{†*}	0.905 ^{†*}	0.834 ^{†*}	0.097 [†]	0.925 [†]	0.695 [†]	0.145 ^{†*}	0.947 [†]	0.686 [†]
MolXPT	350M	0.215 ^{†*}	0.983	0.859 ^{†*}	-	-	-	-	-	-
Text&Chem T5	250M	0.322 ^{†*}	0.943 ^{†*}	0.901 [†]	0.105 [†]	0.849 ^{†*}	0.697 [†]	0.149 [†]	0.898 ^{†*}	0.705
ChatGPT	-	0.139 ^{†*}	0.887 ^{†*}	0.847 ^{†*}	0.044 ^{†*}	0.867 ^{†*}	0.671 ^{†*}	0.048 ^{†*}	0.852 ^{†*}	0.665 ^{†*}
Aug-T5	77M	0.305	0.907	0.877	0.070	0.892	0.700	0.141	0.911	0.685
Aug-T5 _{base}	250M	0.386	0.955	0.884	0.098	0.927	0.696	0.158	0.952	0.681
Ada-T5	77M	0.449	0.967	0.905	0.135	0.945	0.725	0.170	0.955	0.696
Ada-T5 _{base}	250M	0.486	0.974	0.911	0.150	0.956	0.743	0.192	0.969	0.706

Table 4: Results of different models for molecule generation on ChEBI-20, PCdes and DrugBank-23 datasets. [†]/^{*} denotes that Ada-T5_{base}/Aug-T5_{base} perform significantly better than baselines at p -value < 0.01 using t-test. The best scores are in bold. Acc: Accuracy. Val: Validity. MAC: MACCS FTS.

authentic SMILES; but when training on molecule generation, it gets the pseudo description. Consequently, the gap between the input training data leads to the gap between the model performance on different tasks. Furthermore, compared with the results for molecular captioning, the base counterparts of both methods for molecule generation exhibit pronounced enhancements, which could also attributed to the gap between the input data, as using the "pseudo" part as the input for molecule generation might offer more space for improvements, especially for larger-scale models that can better tolerate the "pseudo" data nuances.

The difference between Aug-T5 and Ada-T5 also indicates the importance of data authenticity and the difference between real data and pseudo data: as Ada-T5 is later fine-tuned with 100% real data (in comparison with Aug-T5, which is fine-tuned with the mix of real data and pseudo data), its misunderstandings about SMILES during domain adaptation through pseudo data are corrected and therefore has a better overall performance. This further stresses that using pseudo data for direct application may not be the optimal way to exploit its potential.

Effect of the Amount of Pseudo Data

In order to further demonstrate how the amount of pseudo data could affect model performance, we experiment on ChEBI-20, the largest and most widely used dataset, with varying numbers of pseudo data samples \mathcal{N} from 1k to 512k.

Results on Molecular Captioning Figure 5 shows the results of Ada-T5 and Aug-T5 for molecular captioning with different amounts of pseudo data. Both Ada-T5 and Aug-T5 exhibit significant improvements when a modest amount of pseudo data is incorporated into their training. With just 1k pseudo data, both methods can surpass T5_{large} and ChatGPT and achieve a comparative performance to MolT5_{large} and MolXPT. This phenomena is often seen in other data augmentation strategies (Wei and Zou 2019; Sennrich, Haddow, and Birch 2016), and can be attributed to the moderate noise introduced by the pseudo data, which in turn bolsters model generalization. As the amount of pseudo data increases, Ada-T5 and Aug-T5 exhibit different tendencies. The performance of Aug-T5 begins to decline when the number of pseudo data samples reaches 4k, and sees a sharp drop when it exceeds 32k. This is possibly due to the im-




Figure 5: Results of molecular captioning task using different amount of pseudo data.




Figure 6: Results of molecule generation task using different amount of pseudo data.

balance between real data and pseudo data: As the model becomes increasingly exposed to unreal patterns from the pseudo data, it might shift its attention away from genuine patterns. Consequently, the real patterns are overlooked by the model that focuses on the artificial ones. In contrast, Ada-T5 thrives with the increasing amount of pseudo data, evidenced by the growth of overall metrics. One possible explanation is that Ada-T5 only uses pseudo data for pre-training, with follow-up fine-tuning using real data. Thus, the increase of pseudo data does not twist its grasp of genuine patterns, but instead, further amplifies the proficiency of the model during subsequent training.

Results on Text-Based Molecule Generation Figure 6 shows the results of Ada-T5 and Aug-T5 for molecule generation with different amounts of pseudo data. Ada-T5 shows the same superiority and trend as it does in molecular captioning with more pseudo data incorporated, while Aug-T5 displays a non-linear trend, with the optimal choice of the amount of pseudo data significantly larger than when applying Aug-T5 for molecular captioning. The reason might lie in the dual nature of pseudo data: it introduces both linguistic patterns and noise. Initially, a little bit of pseudo data bolsters model generalization by acting as a regularizer. But as more is added, an overabundance of noise degrades the results. However, once a critical mass of pseudo data is reached, the model starts to recognize more subtle and broader linguistic patterns amidst the noise, which helps in generating more accurate SMILES strings, leading to the observed spike in performance. After this peak, the overwhelming volume of pseudo data might reintroduce the dominance of noise, causing a decrease in performance.

The distinct behavior of Aug-T5 in molecular caption-

ing versus molecule generation highlights their inherent differences. Molecular captioning, being more flexible, can buffer linguistic variations, downplaying minor gains from pseudo data and instead more affected by noise. In contrast, molecule generation requires recognizing specific linguistic cues from descriptions that lead to exact structural changes in the SMILES output, making it more receptive to the subtle intricacies but can also discern and benefit from the subtle patterns present in pseudo data. Overall, these results indicate that the impact of pseudo data varies, depending on its inherent nature and the specific task at hand.

Conclusion

In this paper, we introduce a novel approach that enhances low-resource cross-modal molecule discovery by leveraging artificially-real data generated by LLMs. By incorporating a retrieval-based few-shot prompting strategy, we are able to produce high-quality pseudo molecule-description pairs. To mitigate the scarcity of data, we released two datasets: PseudoMD-1M, the first artificially-real dataset for molecule description, and DrugBank-23, a real molecule-description dataset constructed from a novel source. We propose to use pseudo data for domain adaptation and for data augmentation to explore its optimal utilization. Experiments across different datasets show that the former can best exploit the potential of pseudo data, achieving better performance with less parameters and training data. Furthermore, as the performance of the model continues to benefit from the increasing amount of pseudo data, our approach shows the great potential of pseudo data, thereby providing a novel and promising approach for addressing low-resource challenge in cross-modal molecule discovery.

Acknowledgements

We express our gratitude to the anonymous reviewers for their valuable feedback. This research was supported by the National Key R&D Program of China (2021ZD0113302), the National Natural Science Foundation of China Youth Fund (62206079), and the Heilongjiang Provincial Natural Science Foundation of China (YQ2022F006). We also appreciate Du Xiaoman Technology's support for our research.

References

- Anderson, A. C. 2003. The process of structure-based drug design. *Chemistry & biology*, 10(9): 787–797.
- Bagal, V.; Aggarwal, R.; Vinod, P.; and Priyakumar, U. D. 2021. MolGPT: molecular generation using a transformer-decoder model. *Journal of Chemical Information and Modeling*, 62(9): 2064–2076.
- Banerjee, S.; and Lavie, A. 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.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Christofidellis, D.; Giannone, G.; Born, J.; Winther, O.; Laino, T.; and Manica, M. 2023. Unifying molecular and textual representations via multi-task language modelling. *arXiv preprint arXiv:2301.12586*.
- Curtarolo, S.; Hart, G. L.; Nardelli, M. B.; Mingo, N.; Sanvito, S.; and Levy, O. 2013. The high-throughput highway to computational materials design. *Nature materials*, 12(3): 191–201.
- Cuzzucoli Crucitti, V.; Ilchev, A.; Moore, J. C.; Fowler, H. R.; Dubern, J.-F.; Sanni, O.; Xue, X.; Husband, B. K.; Dundas, A. A.; Smith, S.; et al. 2023. Predictive Molecular Design and Structure–Property Validation of Novel Terpene-Based, Sustainably Sourced Bacterial Biofilm-Resistant Materials. *Biomacromolecules*, 24(2): 576–591.
- Dai, Z.; Zhao, V. Y.; Ma, J.; Luan, Y.; Ni, J.; Lu, J.; Bakalov, A.; Guu, K.; Hall, K.; and Chang, M.-W. 2022. Promptagator: Few-shot Dense Retrieval From 8 Examples. In *The Eleventh International Conference on Learning Representations*.
- Du, Y.; Zhao, S.; Chen, Y.; Bai, R.; Liu, J.; Wu, H.; Wang, H.; and Qin, B. 2023. The CALLA Dataset: Probing LLMs' Interactive Knowledge Acquisition from Chinese Medical Literature. *arXiv preprint arXiv:2309.04198*.
- Durant, J. L.; Leland, B. A.; Henry, D. R.; and Nourse, J. G. 2002. Reoptimization of MDL keys for use in drug discovery. *Journal of chemical information and computer sciences*, 42(6): 1273–1280.
- Edwards, C.; Lai, T.; Ros, K.; Honke, G.; Cho, K.; and Ji, H. 2022. Translation between Molecules and Natural Language. In *2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022*.
- Edwards, C.; Zhai, C.; and Ji, H. 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.
- Ferruz, N.; Schmidt, S.; and Höcker, B. 2022. ProtGPT2 is a deep unsupervised language model for protein design. *Nature communications*, 13(1): 4348.
- Frey, N.; Soklaski, R.; Axelrod, S.; Samsi, S.; Gomez-Bombarelli, R.; Coley, C.; and Gadepally, V. 2022. Neural scaling of deep chemical models.
- Gaudelet, T.; Day, B.; Jamasb, A. R.; Soman, J.; Regep, C.; Liu, G.; Hayter, J. B.; Vickers, R.; Roberts, C.; Tang, J.; et al. 2021. Utilizing graph machine learning within drug discovery and development. *Briefings in bioinformatics*, 22(6): bbab159.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B.; et al. 2023. PubChem 2023 update. *Nucleic acids research*, 51(D1): D1373–D1380.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Li, J.; Liu, Y.; Fan, W.; Wei, X.-Y.; Liu, H.; Tang, J.; and Li, Q. 2023. Empowering Molecule Discovery for Molecule-Caption Translation with Large Language Models: A ChatGPT Perspective. *arXiv preprint arXiv:2306.06615*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, A.; Swayamdipta, S.; Smith, N. A.; and Choi, Y. 2022. WANLI: Worker and AI Collaboration for Natural Language Inference Dataset Creation. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, 6826–6847.
- Liu, Z.; Zhang, W.; Xia, Y.; Wu, L.; Xie, S.; Qin, T.; Zhang, M.; and Liu, T.-Y. 2023. MolXPT: Wrapping Molecules with Text for Generative Pre-training. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 1606–1616. Toronto, Canada: Association for Computational Linguistics.
- Malte, A.; and Ratadiya, P. 2019. Evolution of transfer learning in natural language processing. *arXiv preprint arXiv:1910.07370*.
- Miao, Z.; Li, Y.; Wang, X.; and Tan, W.-C. 2020. Snippet: Semi-supervised opinion mining with augmented data. In *Proceedings of The Web Conference 2020*, 617–628.
- Miller, F. P.; Vandome, A. F.; and McBrewster, J. 2009. Levenshtein distance: Information theory, computer science, string (computer science), string metric, damerau? Levenshtein distance, spell checker, hamming distance.

- Min, S.; Lewis, M.; Zettlemoyer, L.; and Hajishirzi, H. 2022. MetaICL: Learning to Learn In Context. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2791–2809.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 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.
- Patani, G. A.; and LaVoie, E. J. 1996. Bioisosterism: a rational approach in drug design. *Chemical reviews*, 96(8): 3147–3176.
- Polykovskiy, D.; Zhebrak, A.; Sanchez-Lengeling, B.; Golovanov, S.; Tatanov, O.; Belyaev, S.; Kurbanov, R.; Artamonov, A.; Aladinskiy, V.; Veselov, M.; et al. 2020. Molecular sets (MOSES): a benchmarking platform for molecular generation models. *Frontiers in pharmacology*, 11: 565644.
- Preuer, K.; Renz, P.; Unterthiner, T.; Hochreiter, S.; and Klambauer, G. 2018. Fréchet ChemNet distance: a metric for generative models for molecules in drug discovery. *Journal of chemical information and modeling*, 58(9): 1736–1741.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1): 5485–5551.
- Rifaioğlu, A. S.; Atas, H.; Martin, M. J.; Cetin-Atalay, R.; Atalay, V.; and Doğan, T. 2019. Recent applications of deep learning and machine intelligence on in silico drug discovery: methods, tools and databases. *Briefings in bioinformatics*, 20(5): 1878–1912.
- Rogers, D.; and Hahn, M. 2010. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5): 742–754.
- Rubin, O.; Herzig, J.; and Berant, J. 2022. Learning To Retrieve Prompts for In-Context Learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2655–2671.
- Schick, T.; and Schütze, H. 2021. Generating Datasets with Pretrained Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 6943–6951. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Schneider, N.; Sayle, R. A.; and Landrum, G. A. 2015. Get Your Atoms in Order An Open-Source Implementation of a Novel and Robust Molecular Canonicalization Algorithm. *Journal of chemical information and modeling*, 55(10): 2111–2120.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Improving Neural Machine Translation Models with Monolingual Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 86–96.
- Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S. S.; Wei, J.; Chung, H. W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; et al. 2023. Large language models encode clinical knowledge. *Nature*, 1–9.
- Tanimoto, T. T. 1958. Elementary mathematical theory of classification and prediction.
- Wang, H.; Zhao, S.; Qiang, Z.; Li, Z.; Xi, N.; Du, Y.; Cai, M.; Guo, H.; Chen, Y.; Xu, H.; et al. 2023a. Knowledge-tuning Large Language Models with Structured Medical Knowledge Bases for Reliable Response Generation in Chinese. *arXiv preprint arXiv:2309.04175*.
- Wang, W. Y.; and Yang, D. 2015. That’s so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using# petpeeve tweets. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, 2557–2563.
- Wang, Y.; Kordi, Y.; Mishra, S.; Liu, A.; Smith, N. A.; Khashabi, D.; and Hajishirzi, H. 2022. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*.
- Wang, Z.; Liang, L.; Yin, Z.; and Lin, J. 2016. Improving chemical similarity ensemble approach in target prediction. *Journal of cheminformatics*, 8: 1–10.
- Wang, Z.; Liu, T.; Peng, H.; and Fang, Y. 2023b. Advances in molecular design and photophysical engineering of perylene bisimide-containing polyads and multichromophores for film-based fluorescent sensors. *The Journal of Physical Chemistry B*, 127(4): 828–837.
- Wei, J.; and Zou, K. 2019. EDA: Easy Data Augmentation Techniques for Boosting Performance on Text Classification Tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 6382–6388.
- Weininger, D. 1988. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1): 31–36.
- Wishart, D. S.; Feunang, Y. D.; Guo, A. C.; Lo, E. J.; Marcu, A.; Grant, J. R.; Sajed, T.; Johnson, D.; Li, C.; Sayeeda, Z.; et al. 2018. DrugBank 5.0: a major update to the DrugBank database for 2018. *Nucleic acids research*, 46(D1): D1074–D1082.
- Xi, N.; Zhao, S.; Wang, H.; Liu, C.; Qin, B.; and Liu, T. 2023. UniCoRN: Unified Cognitive Signal Reconstruction bridging cognitive signals and human language. *arXiv preprint arXiv:2307.05355*.
- Zeng, Z.; Yao, Y.; Liu, Z.; and Sun, M. 2022. A deep-learning system bridging molecule structure and biomedical text with comprehension comparable to human professionals. *Nature communications*, 13(1): 862.