

LLM-Fusion: A Novel Multimodal Fusion Model for Accelerated Material Discovery

Onur Boyar, Indra Priyadarsini, Seiji Takeda, Lisa Hamada

IBM Research - Tokyo

Onur.Boyar@ibm.com, indra.ipd@ibm.com, seijitkd@jp.ibm.com, Lisa.Hamada@ibm.com

Abstract

Discovering materials with desirable properties in an efficient way remains a significant problem in materials science. Many studies have tackled this problem by using different sets of information available about the materials. Among them, multimodal approaches have been found to be promising because of their ability to combine different sources of information. However, fusion algorithms to date remain simple, lacking a mechanism to provide a rich representation of multiple modalities. This paper presents LLM-Fusion, a novel multimodal fusion model that leverages large language models (LLMs) to integrate diverse representations, such as SMILES, SELFIES, text descriptions, and molecular fingerprints, for accurate property prediction. Our approach introduces a flexible LLM-based architecture that supports multimodal input processing and enables material property prediction with higher accuracy than traditional methods. We validate our model on two datasets across five prediction tasks and demonstrate its effectiveness compared to unimodal and naive concatenation baselines.

Introduction

In recent years, materials science has increasingly integrated AI-driven methods to accelerate the discovery of materials with specific properties. One prominent research direction focuses on the *de novo* generation of materials through generative models combined with optimization techniques (Gómez-Bombarelli et al. 2018; Jin, Barzilay, and Jaakkola 2018; Boyar and Takeuchi 2024). Another significant line of research aims to develop robust property prediction models, essential for screening large libraries of existing materials to identify those that exhibit desired characteristics (Soares et al. 2023a,b; Liu et al. 2023). Both research directions involve unimodal and multimodal approaches. Unimodal approaches use a single representation, such as SMILES strings (Weininger 1988), while multimodal architectures leverage various representations—such as SMILES, SELFIES (Krenn et al. 2020), graphs (Kajino 2019; Kishimoto et al. 2023), fingerprints, and text descriptions. These diverse modalities are embedded and fused into a unified representation, enabling the model to effectively perform target

tasks, whether reconstructing the original inputs or predicting material properties. In this paper, we focus on the latter—building a predictor model using a multimodal model via a novel fusion mechanism.




Figure 1: LLM-Fusion architecture.

A common approach to fusing different modalities is through *naive concatenation* (Hu, Rohrbach, and Darrell 2016; Soares et al. 2023a), which combines the embeddings of each modality into a single, concatenated vector. In property prediction tasks, these concatenated vectors often serve as inputs to predictor models. However, this approach has several limitations. One drawback is that it disregards the relationships among the different modalities. An alternative approach is *contrastive learning*, which aims to structure the latent spaces of each modality such that embeddings from different modalities corresponding to the same material are located in similar regions. While this alignment can be effective, contrastive learning becomes increasingly challenging to scale beyond two modalities due to the exponential growth in the number of modality pairs and the computational complexity of managing numerous positive and negative samples during training (Takeda et al. 2023).

In this paper, we propose LLM-Fusion (see Fig. 1), a novel fusion technique that leverages Large Language Models (LLMs) for multimodal modeling. LLMs are highly effective across diverse domains due to their ability to process complex, unstructured data, and also shown promise in numerical prediction tasks (Vacareanu et al. 2024). Our

approach employs an LLM-based fusion model that efficiently scales with additional modalities, providing fixed-size, information-dense representations of multimodal inputs. This scalability makes LLM-Fusion particularly well-suited for property prediction tasks in materials science. We evaluate LLM-Fusion using combinations of two, three, and four modalities, offering, to our knowledge, the most extensive analysis of performance and scalability among multimodal property prediction models. Our experiments across two datasets and five prediction tasks demonstrate that predictive performance consistently improves as more modalities are included, showcasing the potential of LLM-Fusion.

Related Work

Traditional models in materials science often rely on unimodal representations, such as SMILES strings and molecular fingerprints, for property prediction. For example, ChemBERTa (Chithrananda, Grand, and Ramsundar 2020) applies a BERT-like architecture to SMILES strings (Weininger 1988) for learning powerful representations of materials, which can be used for predicting material properties. Morgan fingerprints (Morgan 1965; Rogers and Hahn 2010), which represent molecules as binary vectors indicating the presence or absence of specific substructures, are also widely used for property prediction. Another representation gaining popularity is SELFIES (Krenn et al. 2020). Studies like Yüksel et al. (2023); Priyadarsini et al. (2024) propose architectures trained on a large corpus of SELFIES representations of materials.

However, multimodal approaches, which integrate multiple types of data, require more general mechanisms and are often inspired by methodologies developed outside the materials science domain. In computer vision, for example, the field of multimodal learning has been pioneered by models such as CLIP (Radford et al. 2021), which employs contrastive learning in text-based image generation. Alternatively, many studies (Hu, Rohrbach, and Darrell 2016; Li et al. 2018; Shi et al. 2018) utilize concatenation of representations from different modalities to perform their target tasks. In the material science community, multimodal approaches are focused on both material generation tasks and property prediction tasks, such as in Liu et al. (2023), where they proposed a methodology for text-based editing of molecules, using a contrastive learning-based approach. Another recent approach is Liu et al. (2024), which uses a cross-attention mechanism to learn the fused representations. In another example, Soares et al. (2023a) utilizes graph-based representation and SMILES representation as two modalities and uses a naive concatenation-based approach for the property prediction task. Similarly, Soares et al. (2023b) uses the SMILES representation and tabular data containing a wide variety of molecular features as two modalities and uses a naive concatenation-based approach for the property prediction task. Among these models, we selected naive concatenation as our main competitor model because, unlike other alternatives, it easily scales beyond two modalities and has proven to be useful in many studies.

Proposed Method

In this section, we first discuss our motivation to use LLM as the fusion model, and then provide the details of the LLM-Fusion architecture.

Motivation

An effective multimodal fusion methodology should be flexible, seamlessly adapting as modalities are added or removed. However, existing studies often lack evaluations of their methods’ adaptability to varying numbers of modalities and do not discuss their applicability in dynamic scenarios. Therefore, there is a need for a fusion model that can adjust to different numbers of modalities and improve its performance as new information becomes available.

LLMs are inherently good at summarization of given contents. They are equipped with powerful self-attention mechanisms that allow them to model complex dependencies. Besides, LLMs are the most popular models of recent literature and there is a rapid development of their capabilities. Motivated by these, we implemented them for a fusion task. By extending their capabilities to model the sequences of modality embeddings, LLM-Fusion leverages the LLM’s ability to capture interactions between different types of modalities. The LLM’s self-attention mechanism dynamically weighs the importance of each modality’s features for the prediction task and fuses multiple modalities into a fixed-length, unified representation. Besides, since our proposed method does not focus on a specific type of LLM, as the capabilities of LLMs increase, their performance on the LLM-Fusion framework can also increase.

LLM-Fusion

Our LLM-Fusion model can work with M modalities. The encoders for each of the M modalities can be of any architecture. LLM-Fusion uses the batch of N 1-dimensional embeddings obtained by these M unimodal encoders. The encoders can be frozen if they are already pretrained, and can be trained from scratch if they are not. In the proposed LLM-Fusion model, we do not fine-tune the pretrained encoders. However, in the case of using pretrained encoders, the embedding size of the pretrained encoder and fusion models input embedding dimension may differ. For this reason, we introduced *optional* projection layers. These are used to match the embedding dimension of each embedding with the input embedding dimension of LLM, d_{LLM} , if needed. In Fig. 1, we assumed that all encoders are pretrained encoders, which may require projection layers to match the embedding dimension of the LLM-Fusion model.

Following the optional projection layer, the encoded vectors or transformed vectors via projection layers are stacked along a new dimension, forming a tensor of shape (N, M, d_{LLM}) . The stacked tensor is enriched with positional encodings that specify the location of each modality within the tensor, providing information to distinguish between different sources of information by treating them as a sequence of modalities. This stacked and positional encoding added tensor serves as input to the LLM, which is fed at the input embeddings layer of the LLM. This strategy allows us to skip the traditional tokenization and positional

encoding addition steps of transformer training. Instead, we provide positional encodings based on modalities and provide the embeddings directly from the embedding layer of the LLM. Using this input, the LLM-Fusion model outputs a fixed-size vector of fused representation used for property prediction, which has a dimension of (N, d_{LLM}) . This vector is obtained by taking the average of the last hidden state of the LLM. Note that, even with the increased number of modalities, the size of the embedding vector remains the same. The final layer of our model is a single linear layer that uses the fused representation as input and provides the property prediction.

Experiments

In this section, we discuss our experimental setup, training details, results, and performance improvement when a larger LLM is used as a fusion model.

Datasets

In our experiments, we used two different datasets. In the first set of experiments, we used the MoleculeNet QM9 dataset (Wu et al. 2018), which is a common benchmark dataset in property prediction tasks. Experiments using this dataset used SMILES, SELFIES, and Morgan fingerprints modalities, and subsets of them. We evaluated our model on HOMO, LUMO, and GAP property prediction tasks.

The second dataset we used was the ChEBI-20 dataset proposed by Edwards et al. (2022), which includes SMILES strings of materials and their text descriptions. The availability of text descriptions enabled us to increase the number of modalities to four, namely SMILES, SELFIES, Morgan fingerprints, and text descriptions. Using this dataset, we trained LLM-Fusion models with different combinations of these modalities and evaluated them in LogP and Quantitative Estimate of Drug-likeness (QED) property prediction tasks.

Fusion Model Selection and Modalities

As the fusion model, we chose GPT-2 Small (Radford et al. 2019), motivated by several factors. With only 114 million parameters, GPT-2 Small is considerably smaller than modern LLMs, making it computationally efficient and allowing for faster training and inference times—an important consideration when dealing with multiple modalities and large datasets. Despite its reduced size, it retains the powerful self-attention mechanisms inherent in LLMs, enabling it to capture complex dependencies and interactions between modalities while keeping computational demands manageable. This balance of efficiency and capability makes GPT-2 Small a practical choice for our fusion model.

In Fig. 2, we provide two example architectures that we used in our experiments. Fig. 2(A) shows the case where three modalities, Morgan fingerprints, SMILES, and SELFIES, are used, where (B) shows the case with four modalities where Morgan fingerprints, SMILES, SELFIES, and Text Descriptions are used. In both models, among the encoders, only the encoder for Morgan fingerprints is trained, while the remaining modality encoders’ weights are kept

frozen. In the case of ChemBERTa used for SMILES modality, its class token is used as the embedding to be provided to LLM-Fusion model, while for SELF-BART (Priyadarsini et al. 2024) used for SELFIES modality, and GPT-2 Small is used for the text description modality, the average of the last hidden state of the respective models are used as the embeddings. For the Morgan fingerprints encoder, we used a single linear layer followed by a ReLU activation function. As demonstrated in Fig. 2, the encoders for Morgan fingerprints, SMILES, and text descriptions encoder provides 768-dimensional embeddings, while SELF-BART provides 1024-dimensional embeddings. Since the embedding size of the LLM-Fusion model, GPT-2 Small, is also 768, we only needed to use a projection layer for the SELFIES modality.

Benchmark Models and Training Details

We compare the performance of LLM-Fusion with unimodal models and naive concatenation-based approaches. Following previous studies (Soares et al. 2023a,b), we used XGBoost (Chen and Guestrin 2016) as the predictor for the unimodal and concatenation cases. The naive concatenation method employs the same modalities as LLM-Fusion to evaluate performance as the number of modalities increases. XGBoost hyperparameters were selected via a random search based on the validation set performance¹. We used predefined training-validation-test splits for both the QM9 (108,446/12,050/13,389 instances) and ChEBI-20 (26,403/3,301/3,299 instances) datasets. LLM-Fusion was trained using the AdamW optimizer with weight decay and PyTorch’s ReduceLROnPlateau learning rate scheduler based on validation loss.

Model	Modalities	HOMO ↓	LUMO ↓	GAP ↓
UM	MFP	0.0076	0.0092	0.0113
UM	SM	0.0102	0.0141	0.0154
UM	SF	0.0160	0.030	0.030
NC	SM+MFP	0.0074	0.0090	0.0113
NC	SF+MFP	0.0076	0.0092	0.0113
NC	SM+SF+MFP	0.0074	0.0090	0.0113
Ours	SM+MFP	0.0063	0.0080	0.0085
Ours	SF+MFP	0.0067	0.0080	0.0079
Ours	SM+SF+MFP	0.0053	0.0055	0.0078

Table 1: Results after fine-tuning the LLM-Fusion (Ours) model and benchmarking against naive concatenation (NC) + XGBoost and unimodal (UM) + XGBoost models.

Results

Tables 1 and 2 summarize the results. In each table, SMILES, Morgan fingerprints, SELFIES modalities denoted as SM, MFP, and SF, respectively. Results are presented as Mean Absolute Error (MAE) in Table 1 and as Mean Squared Error (MSE) in Table 2. As can be seen from the table, we evaluated models with different combinations of modalities, for both our proposal and naive concatenation approach, and it is clear that as the number of modalities

¹Tuned hyperparameters: *learningRate*, *maxDepth*, *nEstimators*, *subSample*, *colSampleByTree*, *gamma*, *minChildWeight*.




Figure 2: LLM-Fusion models used in experiments. **A** shows the three-modality model used for the HOMO, LUMO, and GAP property prediction tasks, **B** shows the four-modality model used for the LogP and QED prediction tasks.

Model	Modalities	LogP ↓	QED ↓
UM	MFP	3.083	0.0091
UM	SM	5.522	0.0111
UM	TEXT	9.131	0.0212
UM	SF	2.521	0.0077
NC	SM+TEXT	4.971	0.0112
NC	SM+TEXT+MFP	2.022	0.0061
NC	SM+TEXT+MFP+SF	2.002	0.0063
Ours	SM+TEXT	2.781	0.0062
Ours	SSM+TEXT+MFP	1.578	0.0034
Ours	SM+TEXT+MFP+SF	1.332	0.0029

Table 2: Results after fine-tuning the LLM-Fusion (Ours) model and benchmarking against naive concatenation (NC) + XGBoost and unimodal (UM) + XGBoost models.

increases, the predictive performance of LLM-Fusion increases. On the other hand, we understand that the naive concatenation method fails to scale its performance with the increased number of modalities. One reason for this phenomenon for the naive concatenation approach is that the dimension of the input vector keeps increasing as new modalities are added, which causes the curse of dimensionality. On the other hand, our approach can enrich the fixed-size fused representation as new information is provided.

A similar pattern is observed in Table 2 for the LogP and QED prediction tasks across the two, three, and four-modality settings. LLM-Fusion consistently outperforms both unimodal models and naive concatenation baselines, achieving lower MSE values for both tasks. At each combination of modalities, LLM-Fusion provides better performance than the naive concatenation method, and its performance improves with the addition of more modalities. This demonstrates that even though the fused representation dimension remains the same, adding new modalities enriches the representation and enhances predictive performance.

Using a Larger LLM

In our experiments, we used the GPT-2 Small model as the fusion model. We conducted a small study to see how the model behaves when it is replaced by a larger model. We re-

placed the GPT-2 Small model with the GPT-2 Large (Radford et al. 2019) model, which has 784 million parameters, and evaluated it using the four-modality case in the ChEBI-20 dataset. GPT-2 Large has an embedding dimension of 1280, therefore, for each modality except Morgan fingerprints, we employed projection layers and trained the LLM-Fusion model for LogP and QED property prediction tasks.

Modalities	LLM Type	LogP ↓	QED ↓
SM+TEXT+MFP+SF	GPT-2 Small	1.332	0.0029
SM+TEXT+MFP+SF	GPT-2 Large	1.270	0.0027

Table 3: Prediction performance when GPT-2 Small model is replaced by GPT-2 Large model.

Table 3 shows our results. For a fair comparison, we trained the GPT-2 Large model using the same GPU time (24 hours) as the GPT-2 Small model, which resulted in a smaller number of training epochs for the GPT-2 Large model due to differences in model sizes. However, the GPT-2 Large model still provides better results compared to the GPT-2 Small model, supporting our claim of obtaining a better fusion model by simply replacing the LLM itself.

Conclusion and Future Work

LLM-Fusion represents a novel approach to multimodal property prediction in material science by leveraging the capabilities of large language models. Our results demonstrate that this architecture can significantly enhance property prediction accuracy, streamline the integration of diverse data types, and provide a scalable solution for material discovery. Future work will explore the application of this model to molecular generation tasks, to be used in scenarios such as text-based editing of materials. Besides, we will also explore the impact of using the more recent family of LLMs such as LLaMa (Touvron et al. 2023) or Granite (Mishra et al. 2024) models as our fusion model. The major drawback of our proposal is the computational complexity of the LLM-Fusion model. For this reason, we also aim to work on a general LLM-Fusion model that is pretrained on a larger number of modalities using a larger corpus of materials, to be readily used in downstream tasks.

References

- Boyar, O.; and Takeuchi, I. 2024. Latent Space Bayesian Optimization With Latent Data Augmentation for Enhanced Exploration. *Neural Computation*, 36(11): 2446–2478.
- Chen, T.; and Guestrin, C. 2016. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. ACM.
- Chithrananda, S.; Grand, G.; and Ramsundar, B. 2020. ChemBERTa: large-scale self-supervised pretraining for molecular property prediction. *arXiv preprint arXiv:2010.09885*.
- Edwards, C.; Lai, T.; Ros, K.; Honke, G.; Cho, K.; and Ji, H. 2022. Translation between molecules and natural language. *arXiv preprint arXiv:2204.11817*.
- Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; and Aspuru-Guzik, A. 2018. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2): 268–276.
- Hu, R.; Rohrbach, M.; and Darrell, T. 2016. Segmentation from natural language expressions. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, 108–124. Springer.
- Jin, W.; Barzilay, R.; and Jaakkola, T. 2018. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, 2323–2332. PMLR.
- Kajino, H. 2019. Molecular hypergraph grammar with its application to molecular optimization. In *International Conference on Machine Learning*, 3183–3191. PMLR.
- Kishimoto, A.; Kajino, H.; Hirose, M.; Fuchiwaki, J.; Priyadarsini, I.; Hamada, L.; Shinohara, H.; Nakano, D.; and Takeda, S. 2023. MHG-GNN: Combination of Molecular Hypergraph Grammar with Graph Neural Network. *arXiv preprint arXiv:2309.16374*.
- Krenn, M.; Häse, F.; Nigam, A.; Friederich, P.; and Aspuru-Guzik, A. 2020. Self-referencing embedded strings (SELF-IES): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4): 045024.
- Li, R.; Li, K.; Kuo, Y.-C.; Shu, M.; Qi, X.; Shen, X.; and Jia, J. 2018. Referring image segmentation via recurrent refinement networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5745–5753.
- Liu, P.; Ren, Y.; Tao, J.; and Ren, Z. 2024. Git-mol: A multi-modal large language model for molecular science with graph, image, and text. *Computers in biology and medicine*, 171: 108073.
- Liu, S.; Nie, W.; Wang, C.; Lu, J.; Qiao, Z.; Liu, L.; Tang, J.; Xiao, C.; and Anandkumar, A. 2023. Multi-modal molecule structure–text model for text-based retrieval and editing. *Nature Machine Intelligence*, 5(12): 1447–1457.
- Mishra, M.; Stallone, M.; Zhang, G.; Shen, Y.; Prasad, A.; Meza Soria, A. L.; Merler, M.; Selvam, P.; Surendran, S.; Singh, S.; Sethi, M.; et al. 2024. Granite Code Models: A Family of Open Foundation Models for Code Intelligence. *arXiv preprint arXiv:2405.04324*.
- Morgan, H. L. 1965. The Generation of a Unique Machine Description for Chemical Structures - A Technique Developed at Chemical Abstracts Service. *Journal of Chemical Documentation*, 5(2): 107–113.
- Priyadarsini, I.; Sharma, V.; Takeda, S.; Kishimoto, A.; Hamada, L.; and Shinohara, H. 2024. Improving Performance Prediction of Electrolyte Formulations with Transformer-based Molecular Representation Model. *arXiv preprint arXiv:2406.19792*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language Models are Unsupervised Multitask Learners. *OpenAI Blog*, 1(8): 9.
- Rogers, D.; and Hahn, M. 2010. Extended-Connectivity Fingerprints. *Journal of Chemical Information and Modeling*, 50(5): 742–754.
- Shi, H.; Li, H.; Meng, F.; and Wu, Q. 2018. Key-word-aware network for referring expression image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 38–54.
- Soares, E.; Brazil, E. V.; Gutierrez, K. F. A.; Cerqueira, R.; Sanders, D.; Schmidt, K.; and Zubarev, D. 2023a. Beyond chemical language: A multimodal approach to enhance molecular property prediction. *arXiv preprint arXiv:2306.14919*.
- Soares, E.; Kishimoto, A.; Brazil, E. V.; Takeda, S.; Kajino, H.; and Cerqueira, R. 2023b. Improving Molecular Properties Prediction Through Latent Space Fusion. *arXiv preprint arXiv:2310.13802*.
- Takeda, S.; Kishimoto, A.; Hamada, L.; Nakano, D.; and Smith, J. R. 2023. Foundation model for material science. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 15376–15383.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971*.
- Vacareanu, R.; Negru, V.-A.; Suci, V.; and Surdeanu, M. 2024. From Words to Numbers: Your Large Language Model Is Secretly A Capable Regressor When Given In-Context Examples. *arXiv preprint arXiv:2404.07544*.
- 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.
- Wu, Z.; Ramsundar, B.; Feinberg, E. N.; Gomes, J.; Geniesse, C.; Pappu, A. S.; Leswing, K.; and Pande, V. 2018. MoleculeNet: A Benchmark for Molecular Machine Learning. *Chemical Science*, 9(2): 513–530.

Yüksel, A.; Ulusoy, E.; Ünlü, A.; and Doğan, T. 2023. SELFormer: molecular representation learning via SELFIES language models. *Machine Learning: Science and Technology*, 4(2): 025035.