

MELT: Materials-aware Continued Pre-training for Language Model Adaptation to Materials Science

Junho Kim^{1*}, Yeachan Kim^{1*}, Jun-Hyung Park², Yerim Oh¹, Suho Kim¹, SangKeun Lee^{1,3}

¹Department of Artificial Intelligence, Korea University, Seoul, Republic of Korea

²Division of Language & AI, Hankuk University of Foreign Studies, Seoul, Republic of Korea

³Department of Computer Science and Engineering, Korea University, Seoul, Republic of Korea

{monocrat, yeachan, yerim0210, suho6347, ya1phy}@korea.ac.kr, jhp@hufs.ac.kr

Abstract

We introduce a novel continued pre-training method, MELT (Materials-aware continued pre-training), specifically designed to efficiently adapt the pre-trained language models (PLMs) for materials science. Unlike previous adaptation strategies that solely focus on constructing domain-specific corpus, MELT comprehensively considers both the corpus and the training strategy, given that materials science corpus has distinct characteristics from other domains. To this end, we first construct a comprehensive materials knowledge base from the scientific corpus by building semantic graphs. Leveraging this extracted knowledge, we integrate a curriculum into the adaptation process that begins with familiar and generalized concepts and progressively moves toward more specialized terms. We conduct extensive experiments across diverse benchmarks to verify the effectiveness and generality of MELT. A comprehensive evaluation convincingly supports the strength of MELT, demonstrating superior performance compared to existing continued pre-training methods. In-depth analysis of MELT also shows that MELT enables PLMs to effectively represent materials entities compared to the existing adaptation methods, thereby highlighting its broad applicability across a wide spectrum of materials science¹.

1 Introduction

Materials science encompasses interdisciplinary studies concerning the behaviors, properties, and applications of materials. Given the vast search space in materials science, deep learning-based approaches have emerged as significant avenues to accelerate the entire research pipeline (Tshitoyan et al., 2019; Weston et al., 2019; Olivetti et al., 2020). Specifically, methods centered on natural

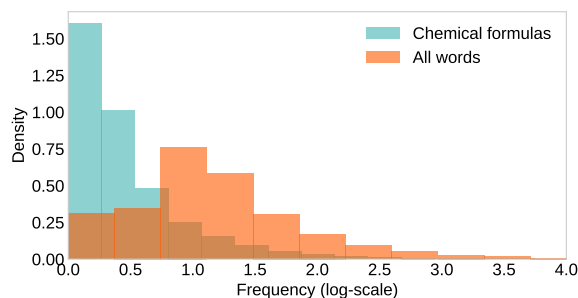


Figure 1: Frequency histograms of all words and chemical formulas on materials science corpus (150K materials-related scientific papers).

language processing (NLP) have provided promising results across a number of materials tasks, such as materials entity recognition (Weston et al., 2019; Friedrich et al., 2020) and synthesis action retrieval (Wang et al., 2022a). However, the limited and scarce nature of datasets in this domain poses substantial challenges for developing models that generalize well across a broad range of materials entities (Song et al., 2023a).

One promising approach to addressing this limitation involves adapting the pre-trained language models (PLMs) for materials science by continuously pre-training them on materials science corpus (Gupta et al., 2022; Huang and Cole, 2022). However, these methods have predominantly focused on constructing the domain-specific *corpora* used in the continued pre-training process, neglecting the *training strategies* employed in the adaptation process. This oversight can lead to inefficiencies in capturing domain-specific nuances and knowledge. For example, the chemical formulas (e.g., LiCoO_2), which are fundamental terms in the materials science field, are typically infrequent words (Figure 1). Therefore, these domain-specific terms are often inadequately captured by the random masking strategy used in previous studies, leading to sub-optimal adaptation results.

* These authors contributed equally to this work.

¹Our code is available at <https://github.com/JunhoKim94/MELT>

In this paper, we propose a novel continued pre-training method, namely MELT (**MatEriaLs**-aware continued pre-Training) that learns the knowledge of materials science through the tailored masking strategy to materials science. To this end, we begin by extracting chemical entities from the materials science corpus. However, these entities alone do not encompass all the fundamental concepts of materials science, specifically the *structure-property-processing-performance* paradigm (William and Callister, 1989). To bridge this gap, we expand the coverage of materials knowledge by constructing semantic graphs that integrate relevant concepts and missing entities. Building on this enhanced materials knowledge, we introduce curriculum into the adaptation process, starting with more familiar and interconnected entities and gradually advancing to more specialized and less common terms. This comprehensive adaptation allows PLMs to effectively learn materials science knowledge.

To verify the efficacy of MELT, we perform extensive experiments, spanning a wide range of downstream tasks for materials science on both generation and classification tasks. Comprehensive evaluation results clearly support the superiority of MELT over strong baselines. Moreover, we verify that MELT enables PLMs to effectively learn the materials by thoroughly transforming both the training corpora and adaptation strategies, underscoring the efficacy of the tailored approach to materials science. In summary, the contributions of this paper include the following:

- We propose MELT, a novel continued pre-training method to adapt the pre-trained language models for materials science.
- We introduce a method for constructing a semantic graph of materials entities to widen the coverage of materials knowledge.
- We demonstrate that MELT substantially improves the performance compared to previous adaptation methods for materials science.

2 Related Works

2.1 NLP in Materials Science

The increasing number of textual datasets in materials science (e.g., scientific publications, patents) has facilitated the use of NLP-based approaches to address various materials tasks, such as relation classification (Mysore et al., 2019a; Mullick

et al., 2024a) and materials entity extraction (Weston et al., 2019; Friedrich et al., 2020). For example, Weston et al. (2019) proposed a bidirectional LSTM tagger for named entity recognition on tags associated with the well-known materials science tetrahedron (i.e., structure, property, processing, performance). Tshitoyan et al. (2019) demonstrated promising results with embedding-based unsupervised methods for understanding chemistry knowledge and chemical properties. Beyond embedding models, Trewartha et al. (2022) introduced PLMs trained on a materials science corpus following the BERT procedure (Devlin et al., 2019). Similarly, Gupta et al. (2022) and Huang and Cole (2022) adapted SciBERT (Beltagy et al., 2019) and BERT (Devlin et al., 2019) to the domains of general materials and battery-specific corpora by additionally training on the domain-specific corpus. Recently, HoneyBee (Song et al., 2023b) suggested the materials domain-specific instruction data to fine-tune the large language models.

Previous studies have demonstrated promising results in adapting PLMs to materials science. However, these methods have primarily focused on adapting the corpus for materials science while employing a basic random masking strategy (Devlin et al., 2019). Such a domain-agnostic approach potentially prevents the models from adequately learning about chemical entities and diverse formulas, which often fall into the less frequent tail distributions in word frequency. In contrast, MELT is a tailored approach to materials science, focusing on both the adaptation corpus and learning strategies to adapt PLMs for materials science.

2.2 Continued Pre-training of PLMs

Building PLMs from scratch requires substantial computational resources; therefore, continued pre-training has garnered significant attention, especially in scientific domains. For example, Gururangan et al. (2020) adapted PLMs trained on general corpora to various domains (e.g., computer science, biomedical) by performing random-based masked language modeling on the target domain corpus. Lin et al. (2021b) performed the domain adaptation by selectively masking entities based on the trained taggers. Subsequently, Wilf et al. (2023) proposed Diff-Masking that utilizes the frequency difference between generic and target domains to perform masked language modeling on domain-specific terms. In addition to the masking strategy, several works studied the catastrophic forgetting

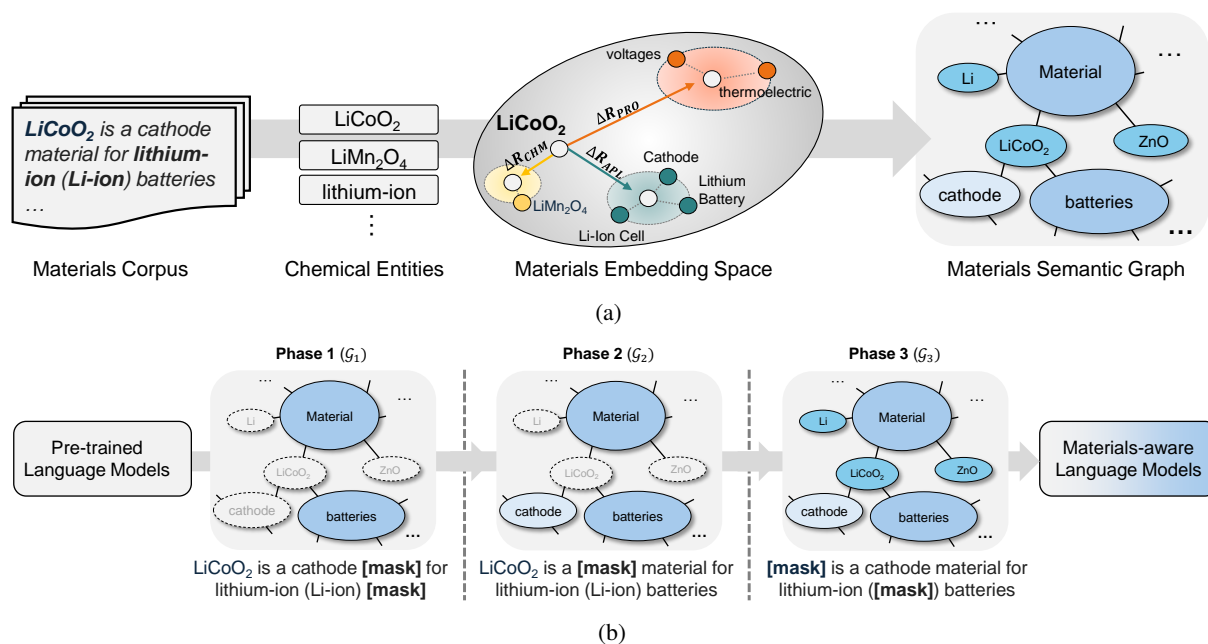


Figure 2: Overall adaptation process of PLMs with MELT. Starting from the materials science corpus, we extract chemical entities (e.g., chemical names and formulas). These entities are then expanded to related terms based on the materials embeddings with their compositional property, resulting in materials semantic graphs. Based on the constructed graphs, MELT performs curriculum masking over materials corpus (e.g., three phases in this example).

problem when adapting PLMs to other domains (Jin et al., 2022; Ke et al., 2023).

In contrast, MELT has distinct characteristics compared to the previous continued pre-training methods. First, previous works have designed a domain-agnostic masking strategy based on random sampling and word frequency. However, MELT involves a tailored masking strategy to materials science by comprehensively constructing materials knowledge based on the fundamental concepts of materials (i.e., structure, property, processing, performance). Second, MELT integrates a curriculum into the adaptation process, providing better local optima compared to naive adaptation methods (Bengio et al., 2009).

3 Materials-aware Continued Pre-training

We elaborate on the proposed continued pre-training method, MELT. The key strategy is to extract the material entities from the scientific papers and inject the extracted knowledge of the materials into the PLMs through a masked language modeling (MLM) objective. To this end, we first extract the materials entities from the materials science corpus. From these entities, we expand the materials knowledge by semantically augmenting the related entities and properties. We then start pre-training on the basic materials entities and pro-

gressively move more specialized knowledge in a curriculum manner. The overall procedures are described in Figure 2.

3.1 Continued Pre-training through MLM

We start by setting the training objective of continued pre-training. Following the promising results of the previous study (Gururangan et al., 2020), we adapt the PLMs to different domains by performing MLM (Devlin et al., 2019) on the domain-specific corpus related to materials science. Specifically, it masks out a small portion of the input sequence (i.e., replace the original word with the special token [mask]) and trains the model to predict the original tokens. Formally, let the set of words to be masked be denoted as \mathcal{G} , the MLM training objective is as follows:

$$\mathcal{L}(x, \hat{x}; \theta) = \sum_{w_m \in \mathcal{G}, w_m \in x} \log P(y = w_m | \hat{x}; \theta) \quad (1)$$

where x and \hat{x} denote the original and masked input sequences based on the set of masking words \mathcal{G} , respectively. In a typical pre-training objective, the set \mathcal{G} is randomly determined without considering the importance of words in the target domain. However, such a random strategy makes the PLMs poorly learn the sparse domain-specialized terms (Wilf et al., 2023), as these terms are rarely selected

for the training. Therefore, the objective of MELT is to fill the masking set \mathcal{G} with the materials knowledge to adapt the PLMs to the domains of materials science effectively.

3.2 Materials-aware Entity Masking

Chemical Entity Extraction While the general domain corpus is widely spread across diverse sources and formats (e.g., Wikipedia, News, SNS), the large-scale corpus for materials science mainly exists in the form of scientific publications (e.g., patents, papers). Moreover, it has several distinct characteristics. First, the materials corpus includes a number of chemical formulas (e.g., H_2O , LiCoO_2), which are crucial components to understanding materials science. For example, in materials benchmarks for token classification, such as named entity recognition, the formulas roughly take up 20% among all entities². Secondly, materials corpus frequently include various representations of domain-specific jargon and abbreviations (e.g., "THF" for *Tetrahydrofuran*, "AZT" for *3'-azido-3'-deoxythymidine*).

To extract these terms to learn materials knowledge, we follow the hybrid approach of ChemDataExtractor (Swain and Cole, 2016; Mavracic et al., 2021), utilizing both dictionary-based mapping and CRF taggers to identify the materials entities. The extracted terms include not only elemental and compound names but also materials characteristics such as density, melting point, and electrical conductivity. This wealth of material-aware entity extraction provides a solid foundation for the continued pre-training.

Semantic Graph for Knowledge Expansion

While the various types of material-aware entities from previous extraction phases involve knowledge about materials, they still lack the fundamental concept of materials science about structure-property-processing-performance paradigm within materials science (William and Callister, 1989).

To augment these fundamental relations, we construct semantic graphs of materials entities by connecting related terms and finding similar yet missing entities with the seed entities. To this end, we first train a lightweight embedding model (Tshityoyan et al., 2019) on the materials corpus³. Based on the learned embeddings of materials terms and entities, we leverage the compositional property

of the embedding representations (Mikolov et al., 2013) where the addition of the two embeddings can indicate the compositional meaning (e.g., the sum of the two embeddings, *Vietnam* and *Capital*, can infer the meaning of *Hanoi*). For example, to find out the *Application* of the *carbon* entity, adding the embedding about *Application* to the embedding of *carbon* can infer new entities related to the property of carbon (Detailed example can be found in Appendix A.1).

To obtain the pivot vectors representing fundamental concepts of materials science, we extract the example pairs of each concept⁴ from MatKG (Venugopal et al., 2022) and build the concept embeddings by subtracting the word embeddings of the pair of each concept. Formally, let the set of the word pairs for the concept R be denoted as \mathcal{S}_R , and the embedding for the concept R is derived as follows:

$$e(R) = \frac{1}{|\mathcal{S}_R|} \sum_{(w_a, w_b) \in \mathcal{S}_R} (e(w_a) - e(w_b)) \quad (2)$$

where w_a and w_b are the subject and object for the relation concept R , and $e(w)$ denotes the trained embedding for the word w . Based on the embeddings of each relation concept, we build the semantic graph where nodes represent entities and edges represent similarity between entities. Specifically, we make the connections based on the cosine similarity between all words in trained embeddings and compositional embedding of extracted entities and concepts as follows:

$$\frac{(e(w_i) + e(R)) \cdot e(\mathcal{V})}{\|(e(w_i) + e(R))\| \|e(\mathcal{V})\|}, \forall w_i \in \mathcal{G} \quad (3)$$

where $e(\mathcal{V})$ denotes the embeddings for all words. However, all words connected by this procedure are not closely related to each entity. Therefore, we only leave the top-k elements based on their similarity. Throughout this process, we obtain a richer set of materials entities for MLM. The number of unique entities is in Appendix A.4

3.3 Curriculum-based Entity Learning

While continued pre-training with the tailored masking can effectively train the knowledge of materials sciences, learning with specialized entities from the beginning makes the training unstable,

²Based on the statistics of MatScholar and SOFC.

³Details for word embeddings are in Appendix A.2

⁴We consider six fundamental concepts: Material, Property, Application, Characterization, Descriptor, Symmetry/Phase. The details for each relation are listed in Appendix A.3.

Algorithm 1 Curriculum-based Entity Learning

Input: Node degree-based curriculum $\{\mathcal{G}_1 \dots \mathcal{G}_K\}$, PLM parameters θ , training corpus D , maximum training step τ , masking probability p_m ,

```
1: for token sequence  $T$  in  $D$  do
2:   Search entities  $e$  in  $T$ .
3:   Order entities according to  $\{\mathcal{G}_1 \dots \mathcal{G}_K\}$ 
4: end for
5: while step  $< \tau$  do
6:   for stage  $i = 1, 2, \dots, K$  do
7:      $T' \leftarrow \text{MaskingEntities}(T, \mathcal{G}_i, p_m)$ 
8:     Train with MLM for PLM  $\theta$  on  $T'$ 
9:   end for
10: end while
11: return PLM parameters  $\theta$ 
```

leading to poor generalization performance (Tay et al., 2019a; Zhan et al., 2021; Lee et al., 2022). Inspired by the human learning process, which learns easy concepts first and moves to harder ones, we introduce a curriculum into the adaptation process.

To define the difficulty within the set of extracted materials terms \mathcal{G} , we utilize the node degree⁵ in the constructed semantic graph in section 3.2. Since the more connected entities are likely to be more familiar and well-defined within the corpus, they can provide a straightforward criterion for prioritizing the learning sequence. Therefore, we gradually mask materials entities from those with higher node degrees (i.e., fundamental and common concepts) to those with lower node degrees (i.e., niche and specialized concepts).

We thus calculate the node degree of all entities and decompose the set of entities \mathcal{G} into a stratified format, i.e., $\mathcal{G} = \{\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_K\}$ where K is the number of stages for the curriculum, and \mathcal{N}_1 consists of the easiest entities while \mathcal{N}_K includes the hardest entities. Then, we gradually expand the materials entities from the easiest entity set \mathcal{G}_1 to the hardest entity set to construct the curriculum set \mathcal{G}_i as follows:

$$\mathcal{G}_i = \mathcal{G}_{i-1} \cup \mathcal{N}_i, \quad (\mathcal{G}_1 = \mathcal{N}_1) \quad (4)$$

By utilizing the curriculum sets, we progressively mask the entities connected to previously learned entities for each stage. The detailed process is represented in Algorithm 1.

⁵Node degree is the number of connected edges to nodes.

4 Experiments

In this section, we verify the efficacy of the MELT. Specifically, we answer the following four questions through extensive experiments and analysis:

Q1 (Adaptability) Does MELT enable better adaptation than existing methods across diverse benchmarks? (§4.2, §4.4 §4.5)

Q2 (Transferability) How the extracted materials knowledge from MELT is related and affect to downstream tasks? (§4.3, §4.6, §4.7)

Q3 (Efficiency) Does MELT offer better efficiency than existing continued pre-training methods for domain adaptation? (§4.8)

Q4 (Insights) What materials science knowledge is extracted and learned during the continued pre-training? (§4.9)

4.1 Experimental Setups

Baselines To confirm the effectiveness of MELT, we mainly compare ours with strong domain adaptation methods for PLMs: **DSP** (Gururangan et al., 2020), **EntityBERT** (Lin et al., 2021a), **Diff-Masking** (Wilf et al., 2023), and **DAS** (Ke et al., 2023). The detailed experimental setups for baselines are described in Appendix C.1.

Pre-training To adapt the PLMs to the domain of materials science, we leverage the 150K scientific papers related to materials science following the previous work (Gupta et al., 2022). In MELT, we implement iterative curriculum learning, consisting of 10K steps for warm-up training and 10K steps for each subsequent curriculum stage. We set the number of curriculum stages (K) to 3 based on the empirical analysis. Detailed analysis is represented in Appendix F.

Downstream Tasks and Datasets To demonstrate the diverse aspects of the MELT, we compare the models on both generation and classification tasks. For generation tasks, we evaluate each baseline using the MatSci-NLP dataset (Song et al., 2023a), which comprises seven materials-related tasks (e.g., materials entity recognition, slot filling). Detailed information on each task within MatSci-NLP can be found in Appendix B. For classification tasks, we adopt four different tasks following the previous research (Gupta et al., 2022), which include NER (MatScholar (Weston et al., 2019), SOFC-EXP (Friedrich et al., 2020)), paragraph

Table 1: Evaluation results on MatSci-NLP. **NER**: Named-Entity Recognition, **RC**: Relation Classification, **EAE**: Event Argument Extraction, **PC**: Paragraph Classification, **SAR**: Synthesis Action Retrieval, **SC**: Sentence Classification, **SF**: Slot Filling. The best and the second best results are highlighted in **boldface** and underline, respectively.

Method	Metric	MatSci-NLP							Overall
		NER	RC	EAE	PC	SAR	SC	SF	
SciBERT (Beltagy et al., 2019)	Micro-F1	0.738	0.818	0.458	0.671	0.701	0.909	0.500	0.693
	Macro-F1	0.517	0.600	0.290	0.568	0.528	0.612	0.258	0.482
DSP (Gururangan et al., 2020)	Micro-F1	0.733	0.857	0.483	0.704	0.772	0.914	0.568	0.722
	Macro-F1	0.500	<u>0.631</u>	<u>0.313</u>	0.579	0.649	0.630	0.322	0.518
EntityBERT (Lin et al., 2021b)	Micro-F1	0.728	0.805	<u>0.495</u>	0.765	0.739	0.897	0.559	0.713
	Macro-F1	0.528	0.567	0.311	0.699	0.621	0.581	0.303	0.516
DAS (Ke et al., 2023)	Micro-F1	0.770	0.848	0.478	0.672	<u>0.778</u>	0.902	<u>0.592</u>	<u>0.725</u>
	Macro-F1	0.567	0.628	0.292	0.607	0.641	0.607	0.356	0.528
Diff-Masking (Wilf et al., 2023)	Micro-F1	<u>0.771</u>	0.858	0.471	0.686	0.769	0.879	0.589	0.721
	Macro-F1	<u>0.575</u>	0.641	0.301	0.573	0.687	<u>0.622</u>	<u>0.373</u>	<u>0.539</u>
MELT (ours)	Micro-F1	0.786	0.860	0.498	<u>0.728</u>	0.798	<u>0.911</u>	0.610	0.741
	Macro-F1	0.593	0.620	0.341	<u>0.647</u>	<u>0.685</u>	0.613	0.395	0.556

Table 2: Ablation results on MatSci-NLP. CEL and MEM indicate curriculum-based entity learning and materials-aware entity masking, respectively. The values in parentheses represent the performance difference to MELT. Detailed results are shown in Appendix E.

Method	MatSci-NLP (Overall)	
	Micro-F1	Macro-F1
MELT (ours)	0.741	0.556
w/o CEL	0.729 (-0.012)	0.542 (-0.014)
w/o CEL, MEM	0.713 (-0.028)	0.516 (-0.040)

classification (Glass Science (Venugopal et al., 2021)), and slot filling (SOFC-Filling (Friedrich et al., 2020)) tasks. Detailed evaluation metrics are shown in Appendix C.3.

Backbones Following prior work (Gupta et al., 2022), we apply each adaptation method to the SciBERT (Beltagy et al., 2019), which is an encoder-based model pre-trained on a scientific corpus comprising approximately 1.14 million documents. For generation tasks, following the recent setup from the previous work (Song et al., 2023a), we construct an encoder-decoder transformer model by integrating a transformer decoder with the encoder model, which pre-trained on each baseline.

4.2 Main Results

Table 1 presents the overall performance results on the MatSci-NLP benchmark. The comparison indicates that MELT yields the best adaptation outcomes compared to previous methods. Specifically, MELT enhances the performance of the backbone model by an average of 6.9% and 15.3% in terms of Micro-F1 and Macro-F1 scores, respectively, across all tasks. This indicates the broad applicability of the proposed method to various materials science tasks. It is also noteworthy that domain-agnostic methods (i.e., Diff-Masking, DAS) reveal limited performance improvement compared to the random masking baseline (i.e., DSP (Gururangan et al., 2020)). It underscores the effectiveness of the domain-specific approach in continued pre-training. Overall, the results strongly support the superiority of the proposed method in adapting PLMs to materials science.

4.3 Ablation Study

We perform ablation studies to confirm whether the components in MELT are crucial in producing better PLMs for materials science. Table 2 shows the ablation results about two components: (i) Materials-aware entity masking (MEM) by building a semantic graph of chemical entities⁶ (ii)

⁶Note that models without the expansion do not have curriculum components as well because the proposed curriculum learning depends on the results of the expansion procedure.

Table 3: Comparison of different curriculum strategies with MELT. The best results are highlighted in **boldface**.

Curriculum Strategy	MatSci-NLP (Overall)	
	Micro-F1	Macro-F1
No Curriculum	0.729	0.542
Frequency	0.736	0.545
Concept	0.736	0.551
Masking Ratio	0.737	0.547
Reverse	0.730	0.552
MELT (ours)	0.741	0.556

Table 4: Evaluation results on four classification tasks. The best and second best results are highlighted in **boldface** and underline, respectively.

Methods	NER _{SOFC}	NER _{MS}	PC	SF
SciBERT	80.3	85.3	94.2	59.1
DSP	80.4	<u>85.8</u>	<u>95.2</u>	61.4
EntityBERT	80.0	85.1	94.7	<u>62.2</u>
DAS	79.7	85.5	95.1	58.4
Diff-Masking	<u>80.6</u>	85.6	94.7	62.0
MELT (ours)	81.1	86.0	95.7	62.9

Curriculum-based entity learning (CEL) based on the node degree of the constructed graph. Notably, we find that omitting each component from MELT leads to a substantial drop in performance, providing empirical verification of the proposed method in continued pre-training. Specifically, we highlight the results of the model without both CEL and MEM in MELT, which uses only the chemical entities similar to those in EntityBERT. Although this model exhibits inferior performance compared to the random masking baseline, it achieves significantly improved performance when expanding its knowledge with fundamental concepts from the materials science tetrahedron and learning it through a well-defined curriculum.

4.4 Effect of the Curriculum Adaptation

To further investigate the effectiveness of the proposed curriculum learning, we compare ours with the existing curriculum-based MLM strategies: Frequency, Concept, Masking Ratio, and Reverse⁷. Table 3 presents the comparison results on the MatSci-NLP benchmark. We observe that the curriculum approach in MELT demonstrates superior perfor-

⁷Details for each strategy is described in the Appendix D

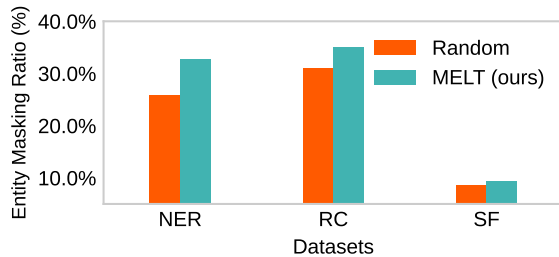


Figure 3: Masking ratio (%) of the materials entities in token classification datasets (NER, RC, SF).

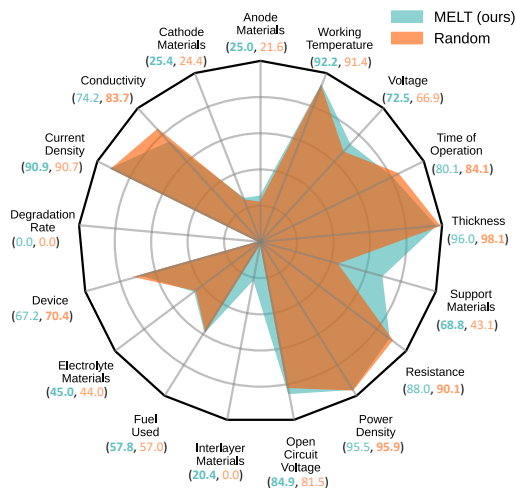


Figure 4: Categorical accuracy (%) for the classes of the chemical entities in the SF task.

mance compared to existing curriculum strategies. These results underscore the efficacy of considering node degree in materials semantic graphs when adapting PLMs for materials science.

4.5 Evaluation on Classification Tasks

To assess the generality of our model, MELT, we evaluate each baseline on four classification tasks following the settings from prior work (Gupta et al., 2022). We present the evaluation results on test sets in Table 4. Similar to the results in generation tasks, MELT outperforms other baselines in most cases. These results demonstrate that PLMs adapted by MELT perform effectively across various classification tasks.

4.6 Masking Relevance to Downstream Tasks

We additionally analyze how the extracted knowledge relates to the entities of the downstream tasks. Here, we compare the proposed method with the random masking. For evaluation, we aggregate all words with their tags in NER (MatScholar and SOFC), RC, and SF tasks and calculate the over-

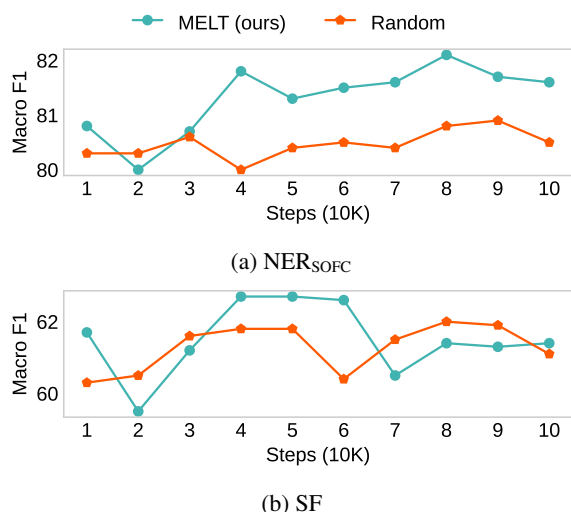


Figure 5: Comparison of the Macro F1 score for the SOFC-NER and SOFC-Filling test sets over the different number of pre-training steps.

lapped ratio of entity tags (labels starting with B- or I-). Figure 3 shows the comparison results with the baseline. The results show that the proposed method, MELT, shows a significantly larger ratio of entity tags. It can be interpreted as our masking strategy successfully extracting domain-specific entities (e.g., chemical formulas, property).

4.7 Effect of Materials-aware Entity Masking

The extracted knowledge from MELT is centered on chemical entities (the first step of the proposed method), and the benchmark datasets for materials science involve a number of chemical entities. To verify the effectiveness of the entity-centric approach, we further analyze the performance of the fine-tuned model for the classes of the chemical entities (i.e., chemical names or formulas) in the slot-filling task (SOFC-Filling). Figure 4 shows the categorical performance of the entity-based (ours) and random-based approaches. We observe that our MELT achieves superior performance in all materials categories. Specifically, MELT outperforms random-based masking about 25% on the Support materials class. Moreover, our MELT performs better than random masking across a variety of properties. These results indicate that materials-aware entity masking can improve the generalization ability of PLMs by learning from diverse material entities.

4.8 Efficiency of Continued Pre-training

We analyze the training efficiency of our proposed MELT by comparing the performance of SOFC-

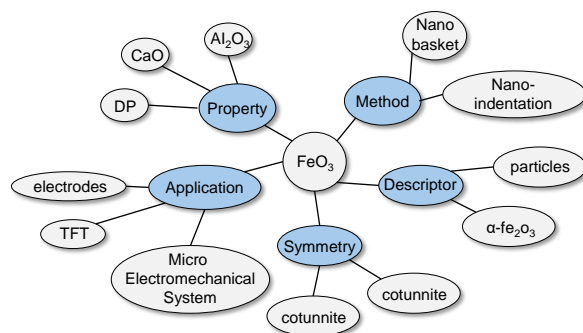


Figure 6: Example of the constructed graphs starting from the initially extracted chemical formula (FeO_3). Only the entities in white circles are used to construct the set of masking words \mathcal{G} .

NER and SOFC-Filling test sets per pre-training step with the baseline. The results are shown in Figure 5. We observe that our MELT consistently outperforms the baselines after 40K steps pre-trained on both datasets. These results indicate that our materials-aware masking and curriculum learning strategies efficiently adapt the PLMs to materials science domains.

4.9 Extracted Materials Knowledge

Our MELT extracts chemical entities from the scientific corpus and expands the knowledge through a materials semantic graph. To validate the constructed knowledge, we sample a part of the graph with the five relation types (i.e., Property, Method, Application, Symmetry label, Descriptor). Figure 6 represents the expansion of material entities (e.g., DP, electrodes, cotunnite) with five relation types from extracted chemical entities in corpora (i.e., FeO_3)⁸. We observe that most extracted entities (e.g., electrodes, Micro Electromechanical Systems) represent the related concepts well, indicating the effectiveness of our automatic knowledge expansion. Moreover, even in cases of entities with incorrect relations (e.g., CaO, Al_2O_3 with the property relation), these entities are still materials-related words, demonstrating that such entities can be effectively used in the training stage. The overall results highlight the usefulness of the extracted knowledge from MELT. More examples can be found in Appendix G

⁸Note that we only represent the additional entities that overlap with MatKG (Venugopal et al., 2022)

5 Conclusion

In this paper, we have proposed MELT, a novel continued pre-training method to adapt PLMs for materials science. Unlike previous approaches that typically focus on the adaptation corpus, our method comprehensively considers both the corpus and the training strategy in the adaptation process. Specifically, we have constructed the materials knowledge base by considering materials science tetrahedron (William and Callister, 1989). This extracted knowledge is then transferred to the PLMs in a structured, curriculum-based manner. We have conducted extensive experiments across diverse materials science benchmarks. The evaluation results convincingly demonstrate that our tailored approach yields superior performance on downstream tasks compared to domain-agnostic methods. In-depth analysis focusing on generality, efficiency, and insights has further supported the efficacy of MELT in continued pre-training, highlighting its broad value in the field of materials science.

Limitations

While we have demonstrated that MELT effectively adapts the PLMs on materials science domains, there are several limitations that present valuable opportunities for future research.

Further Analysis on Generation Tasks We have mainly focused on improving the efficacy of continued pre-training on information retrieval tasks (e.g., NER, Slot Filling), aligning with previous works (Song et al., 2023a; Trewartha et al., 2022; Gupta et al., 2022). However, the applicability of MELT on further generation tasks for material discovery, such as hypothesis or code generation (Miret and Krishnan, 2024), remains under-explored in this work. Nevertheless, considering the significant performance improvements of MELT achieved on generation benchmarks, MatSci-NLP, we believe that MELT is expected to work well within other generation tasks. We leave the exploration of this direction as promising future research.

Tokenization for Materials Sciences We have expanded materials entities and masked them to adapt the PLMs on the materials domain in this work. During this process, we have observed that complex materials terms such as chemical formulas and substances split into multiple subwords (e.g., LiMnO_2 tokenized to {LiMn, ##O, ##2}). Such

over-tokenization may limit the PLMs to learning meaningful representations in the pre-training stage, which leads to sub-optimal results on the downstream tasks (Kononova et al., 2021). These observations highlight the necessity of materials domain-specific tokenizers for processing complex chemical terms, and we expect that future work in this direction will be another promising avenue for domain adaptation for PLMs.

Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.RS-2024-00415812 and No.2021R1A2C3010430) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.RS-2024-00439328, Karma: Towards Knowledge Augmentation for Complex Reasoning (SW Starlab), No.RS-2024-00457882, AI Research Hub Project, and No.RS-2019-II190079, Artificial Intelligence Graduate School Program (Korea University)).

References

- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. Scibert: A pretrained language model for scientific text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*. OpenReview.net.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (NAACL-HLT)*. Association for Computational Linguistics.
- Annemarie Friedrich, Heike Adel, Federico Tomazic, Johannes Hingerl, Renou Benteau, Anika Marusczyk, and Lukas Lange. 2020. The soft-exp corpus and neural approaches to information extraction in the materials science domain. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.
- Tanishq Gupta, Mohd Zaki, NM Anoop Krishnan, and Mausam. 2022. Matscibert: A materials domain

- language model for text mining and information extraction. *npj Computational Materials*, 8(1).
- Suchin Gururangan, Ana Marasovic, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.
- Shu Huang and Jacqueline M. Cole. 2022. Batterybert: A pretrained language model for battery database enhancement. *J. Chem. Inf. Model.*, 62(24).
- Xisen Jin, Dejiao Zhang, Henghui Zhu, Wei Xiao, Shang-Wen Li, Xiaokai Wei, Andrew Arnold, and Xiang Ren. 2022. Lifelong pretraining: Continually adapting language models to emerging corpora. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (NAACL-HLT)*. Association for Computational Linguistics.
- Zixuan Ke, Yijia Shao, Haowei Lin, Tatsuya Konishi, Gyuhak Kim, and Bing Liu. 2023. Continual pre-training of language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*. OpenReview.net.
- Olga Kononova, Tanjin He, Haoyan Huo, Amalie Trewartha, Elsa A Olivetti, and Gerbrand Ceder. 2021. Opportunities and challenges of text mining in materials research. *Iscience*, 24(3).
- Mingyu Lee, Jun-Hyung Park, Junho Kim, Kang-Min Kim, and SangKeun Lee. 2022. Efficient pre-training of masked language model via concept-based curriculum masking. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Chen Lin, Timothy Miller, Dmitriy Dligach, Steven Bethard, and Guergana Savova. 2021a. Entitybert: Entity-centric masking strategy for model pretraining for the clinical domain. In *Proceedings of the Workshop on Biomedical Language Processing*. Association for Computational Linguistics.
- Chen Lin, Timothy A. Miller, Dmitriy Dligach, Steven Bethard, and Guergana Savova. 2021b. Entitybert: Entity-centric masking strategy for model pretraining for the clinical domain. In *Proceedings of the Workshop on Biomedical Language Processing (BioNLP@NAACL-HLT 2021)*. Association for Computational Linguistics.
- Juraj Mavracic, Callum J. Court, Taketomo Isazawa, R. Stephen Elliott, and Jacqueline M. Cole. 2021. Chemdataextractor 2.0: Autopopulated ontologies for materials science. *J. Chem. Inf. Model.*, 61(9).
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, volume 26.
- Santiago Miret and NM Krishnan. 2024. Are llms ready for real-world materials discovery? *arXiv*.
- Ankan Mullick, Akash Ghosh, G Sai Chaitanya, Samir Ghui, Tapas Nayak, Seung-Cheol Lee, Satadeep Bhattacharjee, and Pawan Goyal. 2024a. Matscire: Leveraging pointer networks to automate entity and relation extraction for material science knowledge-base construction. *Computational Materials Science*, 233.
- Ankan Mullick, Akash Ghosh, G Sai Chaitanya, Samir Ghui, Tapas Nayak, Seung-Cheol Lee, Satadeep Bhattacharjee, and Pawan Goyal. 2024b. Matscire: Leveraging pointer networks to automate entity and relation extraction for material science knowledge-base construction. *Computational Materials Science*, 233.
- Sheshera Mysore, Zachary Jensen, Edward Kim, Kevin Huang, Haw-Shiuan Chang, Emma Strubell, Jeffrey Flanigan, Andrew McCallum, and Elsa Olivetti. 2019a. The materials science procedural text corpus: Annotating materials synthesis procedures with shallow semantic structures. In *Proceedings of the 13th Linguistic Annotation Workshop*.
- Sheshera Mysore, Zachary Jensen, Edward Kim, Kevin Huang, Haw-Shiuan Chang, Emma Strubell, Jeffrey Flanigan, Andrew McCallum, and Elsa Olivetti. 2019b. The materials science procedural text corpus: Annotating materials synthesis procedures with shallow semantic structures. In *Proceedings of the Linguistic Annotation Workshop*. Association for Computational Linguistics.
- Elsa A Olivetti, Jacqueline M Cole, Edward Kim, Olga Kononova, Gerbrand Ceder, Thomas Yong-Jin Han, and Anna M Hiszpanski. 2020. Data-driven materials research enabled by natural language processing and information extraction. *Applied Physics Reviews*, 7(4).
- Yu Song, Santiago Miret, and Bang Liu. 2023a. Matscirlp: Evaluating scientific language models on materials science language tasks using text-to-schema modeling. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.
- Yu Song, Santiago Miret, Huan Zhang, and Bang Liu. 2023b. HoneyBee: Progressive instruction finetuning of large language models for materials science. In *Findings of the Association for Computational Linguistics: EMNLP*. Association for Computational Linguistics.
- Matthew C. Swain and Jacqueline M. Cole. 2016. Chemdataextractor: A toolkit for automated extraction of chemical information from the scientific literature. *J. Chem. Inf. Model.*, 56(10).

- Yi Tay, Shuohang Wang, Anh Tuan Luu, Jie Fu, Minh C Phan, Xingdi Yuan, Jinfeng Rao, Siu Cheung Hui, and Aston Zhang. 2019a. Simple and effective curriculum pointer-generator networks for reading comprehension over long narratives. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.
- Yi Tay, Shuohang Wang, Anh Tuan Luu, Jie Fu, Minh C. Phan, Xingdi Yuan, Jinfeng Rao, Siu Cheung Hui, and Aston Zhang. 2019b. Simple and effective curriculum pointer-generator networks for reading comprehension over long narratives. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.
- Amalie Trewartha, Nicholas Walker, Haoyan Huo, Sanghoon Lee, Kevin Cruse, John Dagdelen, Alexander Dunn, Kristin A Persson, Gerbrand Ceder, and Anubhav Jain. 2022. Quantifying the advantage of domain-specific pre-training on named entity recognition tasks in materials science. *Patterns*, 3(4).
- Vahe Tshitoyan, John Dagdelen, Leigh Weston, Alexander Dunn, Ziqin Rong, Olga Kononova, Kristin A Persson, Gerbrand Ceder, and Anubhav Jain. 2019. Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*, 571(7763).
- Vineeth Venugopal, Sumit Pai, and Elsa Olivetti. 2022. The largest knowledge graph in materials science—entities, relations, and link prediction through graph representation learning. In *AI for Accelerated Materials Design NeurIPS 2022 Workshop*.
- Vineeth Venugopal, Sourav Sahoo, Mohd Zaki, Manish Agarwal, Nitya Nand Gosvami, and N. M. Anoop Krishnan. 2021. Looking through glass: Knowledge discovery from materials science literature using natural language processing. *Patterns*, 2(7).
- Zheren Wang, Kevin Cruse, Yuxing Fei, Ann Chia, Yan Zeng, Haoyan Huo, Tanjin He, Bowen Deng, Olga Kononova, and Gerbrand Ceder. 2022a. Ulsa: unified language of synthesis actions for the representation of inorganic synthesis protocols. *Digital Discovery*, 1(3).
- Zheren Wang, Kevin Cruse, Yuxing Fei, Ann Chia, Yan Zeng, Haoyan Huo, Tanjin He, Bowen Deng, Olga Kononova, and Gerbrand Ceder. 2022b. Ulsa: unified language of synthesis actions for the representation of inorganic synthesis protocols. *Digital Discovery*, 1(3).
- Leigh Weston, Vahe Tshitoyan, John Dagdelen, Olga Kononova, Amalie Trewartha, Kristin A Persson, Gerbrand Ceder, and Anubhav Jain. 2019. Named entity recognition and normalization applied to large-scale information extraction from the materials science literature. *Journal of chemical information and modeling*, 59(9).
- Alex Wilf, Syeda Nahida Akter, Leena Mathur, Paul Pu Liang, Sheryl Mathew, Mengrou Shou, Eric Nyberg, and Louis-Philippe Morency. 2023. Difference-masking: Choosing what to mask in continued pre-training. In *Findings of the Association for Computational Linguistics: EMNLP*. Association for Computational Linguistics.
- D William and J Callister. 1989. The structure of crystalline solids. In *Materials Science and Engineering*. Wiley.
- Kyosuke Yamaguchi, Ryoji Asahi, and Yutaka Sasaki. 2020. Sc-comics: A superconductivity corpus for materials informatics. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*. Association for Computational Linguistics.
- Runzhe Zhan, Xuebo Liu, Derek F Wong, and Lidia S Chao. 2021. Meta-curriculum learning for domain adaptation in neural machine translation. In *Association of the Advanced of Artificial Intelligence (AAAI)*, volume 35. AAAI Press.

Appendix

A Details for Entity Expansion

A.1 Example of Chemical Entity Expansion

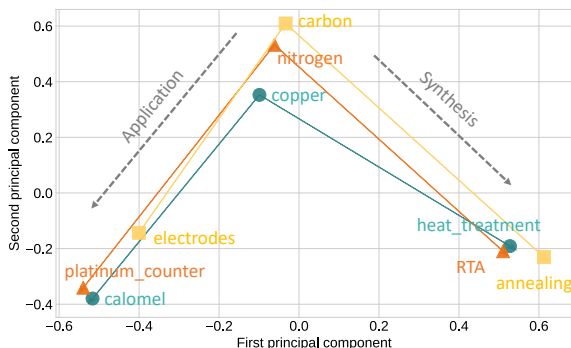


Figure 7: 2D Projection of the materials entities (*carbon*, *nitrogen*, *copper*) having the relations of Applications (*electrodes*, *platinum_counter*, *calomel*) and Synthesis method (*annealing*, *heat_treatment*, *RTA*).

A.2 Details for Word Embedding Models

Table 5: Hyper-parameter settings for our word embedding models

Hyper-parameters	Values
Vocabulary size	876,911
Epochs	30
Embedding sizes	200
Learning rate	0.01
Context Window	8
Sub-sampling threshold	1e-4
Number of negative samples	15

Since the word embedding has been known to capture the relations between words or phrases (Mikolov et al., 2013), we adopt the materials word2vec model, Mat2Vec (Tshitoyan et al., 2019), to expand the materials-related entities and their relations (Section 3.2). In this section, we detail the implementations for word embedding. Following the previous work (Tshitoyan et al., 2019), we tokenized our pre-training corpus, utilizing ChemDataExtractor (Mavracic et al., 2021) and their pre-defined rule-based tokenizer to construct the vocabulary. The vocabulary consisted of all words that occurred more than five times except chemical formulae. We leverage the skip-gram with negative sampling loss. The detailed hyperparameters are shown in Table 5

A.3 Expansion Categories of Materials

We consider six fundamental relations between materials entities, following the relations on the large-scale materials graph, MatKG (Venugopal et al., 2022). Specifically, they categorized text tokens into one of the following seven classes: Material, Property, Application, Synthesis Method, Characterization Method, Descriptor, and Symmetry/Phase Label. Regarding the traditional structure-property-processing-application paradigm in materials science (William and Calister, 1989), these entities encapsulate the total knowledge of any given concept. Thus, we expand the materials properties by utilizing six different categories, which are provided in MatKG (Venugopal et al., 2022).

A.4 Detailed Statistics for Semantic Graph

	#Unique Entities
MELT	713,241
MELT w/o expansion	332,724

Table 6: The number of Unique entities

We construct the materials-aware semantic graph for the tailored masking strategy to materials science (Section 3.2). In this section, we detail the statistics for our semantic graph. Specifically, we compare the number of unique materials-aware entities before and after expansion. The details are shown in Table 6. The results indicate that our expansion methods provide more diverse material-aware entities (e.g., chemical properties and applications).

B Details for MatSci-NLP

Table 7: The detailed size and meta-datasets for MatSci-NLP (Song et al., 2023a).

Task	Size (#Samples)	Meta-Dataset (#Datasets)
NER	112,191	4
RC	25,674	3
EAE	6,566	2
PC	1,500	1
SAR	5,547	1
SC	9,466	1
SF	8,253	1

We compare our models on the multitask generation benchmark, MatSci-NLP (Song et al., 2023a), which includes seven different types of tasks with 13 datasets. Specifically, MatSci-NLP consists of named-entity recognition (NER) (Friedrich et al., 2020; Weston et al., 2019; Yamaguchi et al., 2020; Mysore et al., 2019b), relation classification (RC) (Yamaguchi et al., 2020; Mysore et al., 2019b; Mullick et al., 2024b), event argument extraction (EAE) (Yamaguchi et al., 2020; Mysore et al., 2019b), paragraph classification (PC) (Venugopal et al., 2021), synthesis action retrieval (SAR) (Wang et al., 2022b), sentence classification (SC) (Friedrich et al., 2020), and slot filling (SF) (Friedrich et al., 2020). The detailed statistics are represented in Table 7.

C Implementation Details and Setups

C.1 Continued Pre-training Baselines

We compare our MELT with four strong continued pre-training methods for domain adaptation. For all baselines and MELT, we mask 15% of the total tokens except Diff-Masking. We detail the implementations as follows:

DSP (Gururangan et al., 2020) is a naive approach to continually pre-train the PLMs for domain adaption by utilizing a random masking strategy. We consider the word-level approaches in our experiments following the masking strategies of MatSciBERT (Gupta et al., 2022).

EntityBERT (Lin et al., 2021a) is a method for masking tokens based on whether they are part of domain-specific entities by utilizing the NER model. The original paper leverages the PubMedBERT model, which was trained originally in the clinical domain. In this work, we utilize the ChemDataExtractor (Mavratic et al., 2021) to identify materials entities for masking.

Diff-Masking (Wilf et al., 2023) chooses tokens to mask based on anchor words, which are not commonly found in general domains X_{PT} but commonly found in specific domains X_T . We leverage the X_{PT} as the SciBERT pre-training corpus (Beltagy et al., 2019) and X_T as the corpus, which we constructed in section 3.2. Following the original paper, we mask 25% of the total tokens and use 20 anchor words.

DAS (Ke et al., 2023) is an approach that adaptively updates the parameters by considering the im-

portance of target domains. Following the original paper, we use the weights of contrastive learning λ and temperature τ as 1.0 and 0.05, respectively.

C.2 Hyper-parameters

Pre-train. For hyper-parameters to pre-train the MELT and baselines, we pre-trained all models for 100K steps with a batch size of 128, a sequence length of 128, and a maximum learning rate of $1e-4$. We utilize the linear warm-up ratio 0.048, following the previous work (Gupta et al., 2022). Since the materials entities have different token lengths, we adjust the mask ratios dynamically for each curriculum stage to mask 15% tokens of the token tokens.

Generation Tasks. For generation tasks, we follow the settings in MatSci-NLP (Song et al., 2023a). Specifically, we use a learning rate of $2e-5$, Adam optimizer, batch size of 4, and a maximum number of 20 training epochs with early stopping for MatSci-NLP. We also utilize 3 layers and 8 heads for decoder transformers.

Classification Tasks. For classification tasks, we follow the settings in MatSciBERT (Gupta et al., 2022). Specifically, we use the number of epochs 20, 40, 15, and 10 for SOFC-NER, SOFC-Filling, MatScholar, and Glass Science, respectively. For each task, we sweep the learning rate in $\{2e-5, 3e-5, 5e-5\}$. We use Adam optimizer, and the batch size is 16.

C.3 Evaluation

We follow the evaluation metrics used in MatSci-NLP (Song et al., 2023a) and MatSciBERT (Gupta et al., 2022) for generation and classification tasks, respectively. For the generation benchmark MatSci-NLP, we evaluate each baseline based on both Micro-F1 and Macro-F1. We evaluate each baseline by leveraging both Micro-F1 and Macro-F1 with five different random seeds. For classification, we evaluate each baseline using the Macro-F1 score over entity tags on SOFC-NER and SOFC-Filling, the Micro-F1 score for MatScholar, and the accuracy score for Glass Science tasks. We report the average cross-validation results over five different folds with three random seeds for classification.

D Curriculum Learning Baselines

We compare our models with four different curriculum-based learning methods: frequency,

Table 8: The detailed ablation results for each downstream task on MatSci-NLP. CEL and MEM indicate curriculum-based entity learning and materials-aware entity masking, respectively.

Method	Metric	MatSci-NLP							Overall
		NER	RC	EAE	PC	SAR	SC	SF	
MELT (ours)	Micro-F1	0.786	0.860	0.498	0.728	0.798	0.911	0.610	0.741
	Macro-F1	0.593	0.620	0.341	0.647	0.685	0.613	0.395	0.556
w/o CEL	Micro-F1	0.773	0.830	0.465	0.771	0.778	0.899	0.588	0.729
	Macro-F1	0.565	0.588	0.287	0.683	0.687	0.612	0.373	0.542
w/o CEL, MEM	Micro-F1	0.728	0.805	0.495	0.765	0.739	0.897	0.559	0.713
	Macro-F1	0.528	0.567	0.311	0.699	0.621	0.581	0.303	0.516

masking ratio, and concept-based. We also detail the implementation of baselines. For a fair comparison, all curriculum learning baselines include a warm-up stage. We utilize 3 curriculum stages ($K = 3$) and 10K steps for each curriculum stage.

Frequency. We leverage the frequency of materials entities in the training corpus for defining a difficulty, similar to the previous work (Tay et al., 2019b). We initially train the model with the most frequent materials entities and then progressively add the less frequent entities.

Concept. We utilize the concept-based curriculum learning following the previous work (Lee et al., 2022). We first extract the base concepts using both frequency and node degree and then progressively expand the entities that are related to base concepts.

Masking ratio. We utilize the masking ratio as a difficulty metric, which is used as baselines in previous work (Lee et al., 2022). Specifically, we only mask 10% of the first sequence, and we gradually increase the masking ratio linearly to 20% of tokens when 100K is reached.

Reverse. We train the model with the reverse order of curriculum in MELT. Specifically, we first train with the warmup stage and then progressively train the model from stage 3 to stage 1.

E Detailed Ablation Results

The detailed results are shown in Table 8. We observe that omitting each component in MELT leads to significant performance drops in most downstream tasks.

F The Number of Curriculum Stage

Table 9: Comparison of numbers of curriculum stages on four different classification tasks.

Methods	NER _{SOFC}	NER _{MS}	PC	SF	Overall
2 stage	81.3	86.0	95.4	60.2	80.7
3 stage	81.1	86.0	95.7	62.9	81.4
4 stage	81.0	85.9	96.4	61.9	81.3

To verify the design choice of the number of curriculum stages for MELT, we compare the performance among the model with different numbers of curriculum stages. The results are shown in Table 9. Experimental results show that 3 stage curriculum appears to be best. Based on these results, we use 3 stages ($K = 3$) in all our models.

G Extracted Materials Knowledge

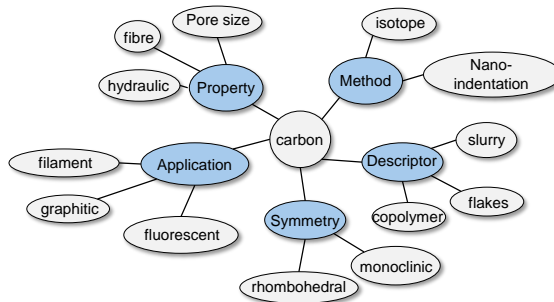


Figure 8: Examples of the constructed graphs starting from the two different chemical entities- Chemical names (*carbon*). Only the entities in white circles are used to construct the set of masking words \mathcal{G} .