TRANSFORMER AS A NEURAL KNOWLEDGE GRAPH

Yuki Nishihori[1] **Yusei Ito**[1] **Yuta Suzuki**[2] **Ryo Igarashi**[3] **Yoshitaka Ushiku**[3] **Kanta Ono**[1]
[1]Osaka University [2]Toyota Motor Corporation [3]OMRON SINIC X Corporation

ABSTRACT

In this study, we propose an effective contrastive learning method that bridges crystal structures with their linguistic properties (*e.g.*, superconductor). Contrastive learning enables both the retrieval of crystal structures based on linguistic characteristics and the inference of linguistic properties from crystal structures, which are essential for accelerating materials discovery. However, a major challenge lies in the limitation of available datasets, which currently include only crystal structures paired with their corresponding article titles and abstracts. Because many papers depend on referenced works and shared domain knowledge—often explored in detail within the main text—titles and abstracts alone do not sufficiently capture the full characteristics of a crystal. To address this issue, we introduce a *neural knowledge graph* by incorporating a transformer into the text encoder of the existing contrastive learning framework, rather than expanding the dataset. This modification enables the model to dynamically incorporate related knowledge, thereby enhancing its representation of linguistic properties and facilitating more accurate correlations between crystal structures and their properties.

1 INTRODUCTION

The crystal structure and physical properties are closely related, making it essential to understand this relationship for material discovery (Callister, 2006). Traditionally, materials modeling has focused on predicting precisely quantifiable properties (*e.g.*, bandgap, formation energy) (Xie & Grossman, 2018; Chen & Ong, 2022; Ito et al., 2025). However, materials discovery is not only driven by such precise numerical targets. Qualitative, linguistically defined properties (*e.g.*, superconductor, ferromagnetic), often offer more practical guidance.

Suzuki et al. (2025) introduced Contrastive Language–Structure Pre-training (CLaSP) to tackle this issue. Their approach utilizes a CLIP-like contrastive learning strategy (Radford et al., 2021), incorporating a text encoder to capture linguistic properties and a crystal encoder to represent crystal structures, thereby linking crystal structures to physical property keywords extracted from academic papers.

However, a major challenge remains: *the lack of sufficiently annotated data for effective training*. Existing datasets, such as the Crystallography Open Database (COD) (Gražulis et al., 2009) and the Inorganic Crystal Structure Database (ICSD) (Belsky et al., 2002), pair crystal structures with paper titles and abstracts. Since research papers typically build upon prior studies or established knowledge, offering detailed analysis within the main text, titles and abstracts rarely provide a complete representation of a crystal structure.

Although knowledge graphs have been proposed as datasets for incorporating contextual knowledge (Venugopal & Olivetti, 2024; Zhang et al., 2024), their direct integration into contrastive learning models remains challenging. A key limitation is that keywords in articles are often missing as entities in existing knowledge graphs. This mismatch between datasets and knowledge graphs raises a new idea: *by representing a knowledge graph within a neural network framework and jointly training it with contrastive learning, we can effectively incorporate peripheral knowledge*.

In this paper, we introduce a *neural knowledge graph* which is a neural network representation of a knowledge graph. This involves a simple modification—adding a transformer (Vaswani et al., 2017) after the original CLaSP's text encoder— to allow the model to dynamically incorporate and learn related knowledge during the contrastive learning process (Fig. 1).

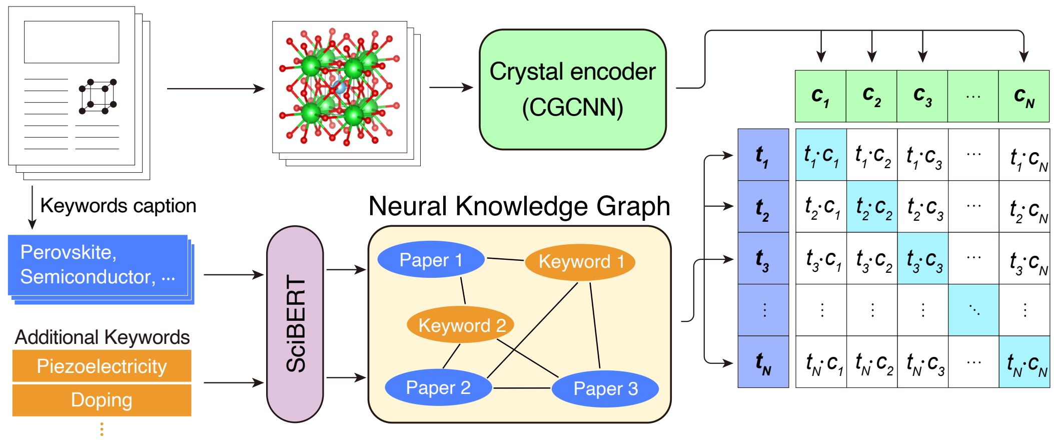

Figure 1: Overall architecture of CLaSP with neural knowledge graph.

2 Neural knowledge graph

In this section, we introduce the transformer as a neural network-based representation of a knowledge graph, called Neural Knowledge Graph (NKG), and describe its integration into CLaSP, a contrastive learning framework.

2.1 Representing a knowledge graph with transformer

NKG utilizes the transformer architecture to process text embeddings encoded from keywords in the titles and abstracts of articles. It dynamically retrieves and integrates relevant knowledge by leveraging the attention mechanism within the transformer, enabling a more contextualized representation of linguistic information.

The simplified one layer of transformer architecture can be described as

$$\text{Transformer}(X, Y) = X + \text{Attention}(XW_q, YW_k, YW_v), \tag{1}$$

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V, \tag{2}$$

where $X \in \mathbb{R}^{N\times d}$ and $Y \in \mathbb{R}^{M\times d}$ are input d-dimensional embeddings with N and M elements, and W_q, W_k, and $W_v \in \mathbb{R}^{d\times d}$ are learnable matrices that project the input embeddings into query, key, and value spaces, respectively. In essence, the attention mechanism calculates the similarity between each element in X (queries) and each element in Y (keys) through a scaled dot product. The resulting similarity scores are then used to weight the corresponding elements in Y (values), which are subsequently added to X. This implies that the feature vector X is updated by dynamically incorporating contextually relevant knowledge from Y.

Applying this to our problem, we aim to incorporate relevant information into the keyword embeddings extracted from each paper. For a collection of N papers, we generate a keyword embedding $\boldsymbol{x_i} \in \mathbb{R}^d$ $(i = 1, 2, \ldots, N)$ for each paper using a text encoder (described later). We also prepare M embeddings $\boldsymbol{y_i} \in \mathbb{R}^d$ $(i = 1, 2, \ldots, M)$ for potentially related keywords. Defining $X = [\boldsymbol{x_1}, \boldsymbol{x_2}, \ldots, \boldsymbol{x_N}]^T$ and $Y = [\boldsymbol{y_1}, \boldsymbol{y_2}, \ldots, \boldsymbol{y_M}]^T$, the transformer, via its attention mechanism, computes the relationships between these embeddings and incorporates the highly related keywords selected from Y into the embedding of each paper. This process can be interpreted as constructing and utilizing a knowledge graph, where the extracted keywords from the papers and the predefined related keywords are treated as entities. We propose the following two methodologies to prepare the set of potentially relevant keywords Y.

Self-attention. The first method is to set $Y = X$, which means using self-attention. This allows the model to integrate information from other papers within the dataset into the representation of each paper. For example, if one paper mentions "superconductivity" and another mentions "high-temperature superconductivity", self-attention can help the model learn a richer representation of "superconductivity" by incorporating information from the related term "high-temperature".

Cross-attention. Another approach is to prepare a set of highly relevant keywords in advance for Y. Unlike the self-attention method, this approach allows updating the embedding by referring to information that is often not explicitly stated in the title and abstract in the paper (*e.g.*, commonly known knowledge, in-depth analysis from the main text or related works). For instance, if a paper mentions "perovskite," the cross-attention mechanism can incorporate information from predefined keywords like "oxide," "bandgap" and "photovoltaics." The predetermined set of keywords are expected to be selected based on expert knowledge to ensure their high relevance. However, if relevance can be sufficiently guaranteed, it is also acceptable to have an LLM generate the keyword set by tuning the prompt or by performing a manual check.

Note that although traditional knowledge graphs typically feature an explicit hierarchical structure and ontology (Hogan et al., 2021), the transformer-based implementation here uses fully connected attention and does not include them. Based on its ability to incorporate peripheral knowledge, we call this a "knowledge graph." In Sec. 4, we discuss in detail the comparison between the traditional knowledge graph and the proposed NKG.

2.2 TEXT ENCODER WITH NEURAL KNOWLEDGE GRAPH

To perform the contrastive learning with our proposed NKG, we incorporated it into CLaSP. CLaSP comprises a crystal encoder, which encodes crystal structures, and a text encoder that encodes keywords extracted from the titles and abstracts of the corresponding papers. These keywords are generated using Meta's Llama3.1 (8B Instruct) (AI@Meta, 2024), which is provided with each paper's title and abstract and instructed to extract ten keywords based on a predefined prompt detailed in the Appendix A. The overall architecture remains consistent with the original CLaSP, with the only modification being the integration of NKG into the text encoder. The model is trained using mini-batches, refining embeddings to draw each sample closer to itself while pushing it away from other samples within the batch. We first explain the text encoder incorporating NKG, followed by a brief overview of the other architecture.

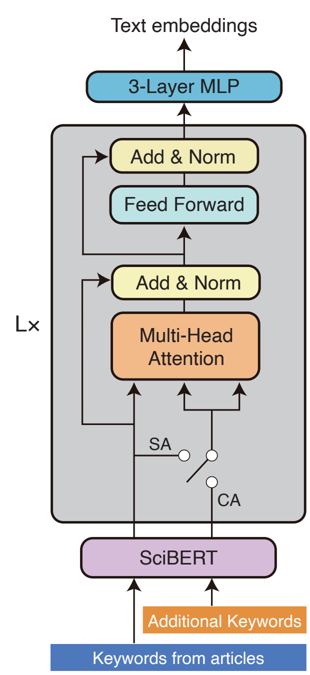

Figure 2: Architecture of text encoder. (SA:Self-Attention, CA:Cross-Attention. In blend-attention, SA and CA switch alternately.)

Text encoder. The extracted keywords from the paper are processed using the frozen pretrained SciBERT model (Beltagy et al., 2019), where they are converted into embeddings via the `CLS` token. This step yields N text embeddings, denoted as $X \in \mathbb{R}^{N \times d}$, where N represents the batch size and d is the embedding dimension. For cross-attention, M keywords that are highly likely to be related to the keywords extracted from the paper were prepared and converted into embeddings using the same SciBERT model, say $Y \in \mathbb{R}^{M \times d}$. To introduce NKG, we employed the original transformer (Vaswani et al., 2017), which consists of multi-head attention, feed-forward networks, residual connections, and layer normalization (see Fig. 2). The only difference is that, in cross-attention method, X serves as the query, while Y is used as both the key and the value. After the transformer layer, the final text embeddings $\boldsymbol{t_i}$ $(i = 1, 2, ..., N)$ is obtained through a three-layer multilayer perceptron (MLP).

Other architecture. The crystal structures are converted into d-dimensional embeddings $\boldsymbol{c_i}$ ($i = 1, 2, ..., N$) by using CGCNN (not pretrained) Xie & Grossman (2018). Contrastive learning is conducted by minimizing the large margin cosine loss function (Wang et al., 2018), as shown in the following equation, to train the model for increasing the cosine similarity of positive pairs while reducing that of negative pairs.

$$L = -\frac{1}{N}\sum_{i=1}^{N}\log\left(\frac{\exp\big(s\big(\cos(\boldsymbol{c_i},\boldsymbol{t_i})-m\big)\big)}{\exp\big(s\big(\cos(\boldsymbol{c_i},\boldsymbol{t_i})-m\big)\big)\;+\;\sum_{j=1,j\neq i}^{N}\exp\big(s\,\cos(\boldsymbol{c_i},\boldsymbol{t_j})\big)}\right), \tag{3}$$

where $s > 0$ is a scaling hyperparameter and $m \in [0, 1]$ is a margin hyperparameter.

3 Experiment

We collected 108,795 crystal structures and their corresponding papers from the COD, dividing the dataset into training data, validation data, and test data in a ratio of 8:1:1. For the cross-attention method, we used an LLM to generate candidate keywords related to physical properties (*e.g.*, Hall effect) and potential causal factors (*e.g.*, Effective mass). We then manually examined these candidates and carefully selected 256 keywords to ensure high relevance and precision. The full list of keywords is provided in the Appendix B.

To evaluate the methods of NKG, we conducted experiments with original CLaSP (without NKG) and three different NKG methods: a 4-layer transformer with all self-attention, a 4-layer transformer with all cross-attention, and a 4-layer transformer alternating between self and cross attention layers (blend-attention). All attention calculations were performed using multi-head attention with 8 heads. The embedding dimensionality for both modalities was set to 768. We optimized the loss function shown in Eq. 3 with a scaling factor s of 3 and a margin m of 0.5, which are the best parameters of the original CLaSP. We used stochastic gradient descent with a batch size N of 1024, trained on four NVIDIA V100 GPUs (resulting in a global batch size of 4×1024). We employed the AdamW optimizer (Loshchilov & Hutter, 2019) with a constant learning rate of 1×10^{-6}, without any warm-up period, and trained for 2000 epochs.

We evaluated the ability to retrieve crystal structures based on a keyword of material properties (*e.g.*, Ferromagnetic) to assess the capability of linking crystal structures with linguistic features. We used cosine similarity to match keyword embeddings with structure embeddings from the test set. We considered a structure to have a specific property if the keywords generated from the corresponding paper's title and abstract contained the target keyword or its variations. We assessed the performance using ROC-AUC metrics which is the area under the ROC (receiver operating characteristic) curve to analyze the trade-off between true and false positives.

Table 1 presents the ROC-AUC results for four keywords, comparing them against the original CLaSP (baseline) and our proposed Neural Knowledge Graph, which incorporates self-attention, cross-attention, and a blended approach. The corresponding ROC curves for each keyword can be found in Appendix C. For three of the four keywords, cross-attention exhibited superior performance, which can be attributed to the fact that the predefined keywords used in cross-attention provide additional context and relevant information that are not explicitly stated in the title or abstract of a typical paper, but are crucial for understanding the properties of the material. Additionally, self-attention performed better than the baseline method for two keywords, likely due to its ability to integrate information from related studies not explicitly mentioned in the paper itself.

Table 1: ROC-AUC comparison for keyword-based crystal structure retrieval tasks. **Bold** indicates the best results, underline the second best.

	Ferromagnetic	Ferroelectric	Semiconductor	Electroluminescence
CLaSP (baseline)	**0.749**	0.686	0.553	<u>0.865</u>
— w/ NKG (self-attention)	<u>0.686</u>	0.688	0.718	0.639
— w/ NKG (cross-attention)	0.669	**0.787**	**0.848**	**0.900**
— w/ NKG (blend-attention)	0.482	<u>0.725</u>	<u>0.812</u>	0.793

4 DISCUSSION AND LIMITATIONS

Keywords for cross-attention method. Although the cross-attention method outperforms the baseline in terms of ROC-AUC, it likely remains reliant on the chosen keyword list. Notably, performance degradation was observed compared to the baseline when evaluated using the keyword "Ferromagnetic". In this preliminary study, the keyword list was constructed through a provisional approach. This limitation suggests that a more systematic and comprehensive approach to keyword selection, incorporating multiple perspectives and keywords strongly associated with physical properties, could further enhance performance.

Comparison with traditional knowledge graphs. One major difference between traditional Knowledge Graphs (KG) and NKG is whether they include an explicit hierarchical structure and ontology. The transformer in NKG can learn these relationships through training, yet it is believed that providing the structure explicitly may affect performance. This raises an interesting question: which approach yields better results—the non-fully connected graph neural network structure with an explicitly defined hierarchical structure (which constrains representational capacity in advance) (Schlichtkrull et al., 2018) or the fully connected attention used in this study? Our transformer model has the potential to generate physically impossible associations, while at the same time revealing important relationships that researchers might have otherwise overlooked.

Direct evaluation of the obtained NKG. In the proposed method, a knowledge graph is represented as a transformer and trained during the contrastive learning process. Therefore, visualizing the transformer model (*i.e.*, displaying the learned knowledge graph) not only allows us to verify whether it aligns with the researchers' intuition but also offers the potential to extract valuable insights in materials science. Various methods for visualizing transformers have been proposed, and applying these diverse approaches is an important future direction (Binder et al., 2016; Selvaraju et al., 2017; Chefer et al., 2021).

5 CONCLUSION

In this work, we presented NKG, a neural network representation of a knowledge graph that leverages a transformer architecture for contrastive learning between crystal structures and their linguistic properties. This approach addresses the common limitations of existing databases, where only paper titles and abstracts are available—often insufficient for comprehensive representation. We demonstrated NKG's effectiveness in a crystal structure retrieval task based on physical property keywords, showing that it outperforms a standard baseline in terms of ROC-AUC. These results highlight NKG's ability to meaningfully integrate relevant knowledge, suggesting that it can enhance the linguistic understanding of crystal structures and potentially accelerate material discovery.

ACKNOWLEDGMENTS

This work is partly supported by the JST-Mirai Program (grant number JPMJMI19G1), the MEXT Program: Data Creation and Utilization-Type Material Research and Development Project (Digital Transformation Initiative Center for Magnetic Materials) (grant number JPMXP1122715503), the MEXT Program: Developing a Research Data Ecosystem for the Promotion of Data-Driven Science, and the JST Moonshot R&D (grant number JPMJMS2236).

The authors thank Tatsunori Taniai, Naoya Chiba, Kotaro Saito, and the anonymous reviewers for their valuable feedback on this study.

REFERENCES

AI@Meta. Llama 3 Model Card. 2024.

Alec Belsky, Mariette Hellenbrandt, Vicky Lynn Karen, and Peter Luksch. New developments in the Inorganic Crystal Structure Database (ICSD): accessibility in support of materials research and design. *Acta Crystallogr. B*, 58(3 Part 1):364–369, 2002.

Iz Beltagy, Kyle Lo, and Arman Cohan. SciBERT: Pretrained Language Model for Scientific Text. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP 2019)*, 2019.

Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, Klaus-Robert Müller, and Wojciech Samek. Layer-Wise Relevance Propagation for Neural Networks with Local Renormalization Layers. In *the 25th International Conference on Artificial Neural Networks (ICANN 2016)*, pp. 63–71, 2016.

W.D. Callister. *Materials Science and Engineering: An Introduction 7th Edition with Wiley Plus Set*. Wiley Plus Products Series. John Wiley & Sons, Incorporated, 2006.

Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2021)*, pp. 782–791. IEEE, 2021.

Chi Chen and Shyue Ping Ong. A universal graph deep learning interatomic potential for the periodic table. *Nat. Comput. Sci.*, 2(11):718–728, 2022.

Saulius Gražulis, Daniel Chateigner, Robert T. Downs, A. F. T. Yokochi, Miguel Quirós, Luca Lutterotti, Elena Manakova, Justas Butkus, Peter Moeck, and Armel Le Bail. Crystallography Open Database – an open-access collection of crystal structures. *J. Appl. Crystallogr.*, 42(4):726–729, 2009.

Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia D'amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. Knowledge Graphs. *ACM Comput. Surv.*, 54 (4):71, 2021.

Yusei Ito, Tatsunori Taniai, Ryo Igarashi, Yoshitaka Ushiku, and Kanta Ono. Rethinking the role of frames for se(3)-invariant crystal structure modeling. In *the 13th International Conference on Learning Representations (ICLR 2025)*, 2025.

Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *the 7th International Conference on Learning Representations (ICLR 2019)*, 2019.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *the 28th International Conference on Machine Learning (ICML 2021)*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021.

Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling Relational Data with Graph Convolutional Networks. In *the 15th Extended Semantic Web Conference (ESWC 2018)*, pp. 593–607, 2018.

Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *16th IEEE International Conference on Computer Vision (ICCV 2017)*, pp. 618–626, 2017.

Yuta Suzuki, Tatsunori Taniai, Ryo Igarashi, Kotaro Saito, Naoya Chiba, Yoshitaka Ushiku, and Kanta Ono. Contrastive Language-Structure Pre-training Driven by Materials Science Literature, 2025. arXiv:2501.12919.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS 2017)*, volume 30, 2017.

Vineeth Venugopal and Elsa Olivetti. MatKG: An autonomously generated knowledge graph in Material Science. *Sci. Data*, 11(217), 2024.

Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. CosFace: Large Margin Cosine Loss for Deep Face Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2018)*, pp. 5265–5274. IEEE, 2018.

Tian Xie and Jeffrey C. Grossman. Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties. *Phys. Rev. Lett.*, 120(14):145301, 2018.

Yuwei Zhang, Fangyi Chen, Zeyi Liu, Yunzhuo Ju, Dongliang Cui, Jinyi Zhu, Xue Jiang, Xi Guo, Jie He, Lei Zhang, Xiaotong Zhang, and Yanjing Su. A materials terminology knowledge graph automatically constructed from text corpus. *Sci. Data*, 11(600), 2024.

A PROMPT FOR EXTRACTING KEYWORDS FROM EACH PAPER

We generated 10 keywords from each paper using Llama with the prompt shown below. The prompt was carefully designed so that it does not generate keywords unrelated to material properties, such as general methods or measurement techniques.

```
def prompt_format_func(material_id, title, abstract):
    """Formats the prompt for the Llama model."""
    prompt_template = """Below are title-abstract pairs for materials
        science papers dealing with crystal structures. For each paper,
         list up to 10 keywords in English that describe the features,
        functions, or applications of the material discussed. Focus on
        the material itself, and do not include general terms or
        measurement techniques (e.g., Crystal Structure, Crystal
        Lattice, X-ray diffraction, Neutron Diffraction, Powder
        Diffraction). Return the results in json format with the
        following schema.

    **Example input 1:**

    ```
    ID: 0001
    Title: Enhancement of Critical Temperature in Layered Copper Oxide
        Superconductors via Lattice Compression Techniques
    Abstract: Superconductivity in copper oxides (cuprates) offers vast
         potential for technological applications due to their high
        critical temperatures (Tc). Our research presents a novel
        approach to enhance Tc in layered cuprate materials through the
         controlled application of lattice compression. Using advanced
        crystallographic methods, we systematically altered the
        interlayer spacing and analyzed the resultant changes in
        electronic properties. Our findings demonstrate a significant
        improvement in superconducting behavior at elevated
        temperatures, further supporting the unconventional mechanisms
        underpinning superconductivity in these materials.
    ```

    Example output 1:

    ```json
    [  {
        "ID": "0001",
        "Keywords": [
          "High-Tc",
          "Cuprate Superconductors",
          "Lattice Compression",
          "Electronic Properties",
          "Layered Structures",
          "Superconducting Phase",
          "Temperature Enhancement",
          "Unconventional Superconductivity"
```
```

```
      ]
    }]
  ```

  Example input 2:

  ```
  ID: 0002
  Title: Advancements in Biodegradable Polymers for Sustained Drug
      Delivery Systems
  Abstract: The development of biocompatible and biodegradable
      materials is critical in the field of medical implants and drug
       delivery systems. This paper examines the latest advancements
      in biodegradable polymers tailored for sustained release of
      therapeutic agents. We analyze various polymer compositions
      that provide controlled degradation rates and compatibility
      with a range of drugs. Our results show promising applications
      in long-term treatments, reducing the need for repeated
      administration and improving patient compliance.
  ```

  Example output 2:

  ```json
  [  {
      "ID": "0002",
      "Keywords": [
        "Biomaterials",
        "Biodegradable Polymers",
        "Sustained Release",
        "Drug Delivery Systems",
        "Biocompatibility",
        "Controlled Degradation",
        "Therapeutic Agents",
        "Medical Implants",
        "Long-Term Treatment"
      ]
    }]
  ```
  """
  prompt = prompt_template + f"""
  Input :

  ```
  ID: {material_id}
  Title: {title}
  Abstract: {abstract}
  ```

  Output :

  ```json
  """
  return prompt
```

B KEYWORDS FOR CROSS-ATTENTION

We used the following 256 keywords as additional keywords for cross-attention method. These keywords were generated using an LLM and were then manually verified to ensure that they are related to physical properties.

```
- Physical Properties                  - Superconductor
- Dielectric                           - Magnetic susceptibility
```

```
- Thermal conductivity
- Ferromagnetism
- Diamagnetism
- Pyroelectricity
- Curie temperature
- Band gap
- Density of states
- Capacitance
- Refractive index
- Optical conductivity
- Electron mobility
- Permeability
- Conductance
- Bulk modulus
- Poisson's ratio
- Elasticity
- Fracture toughness
- Specific heat capacity
- Melting point
- Vapour pressure
- Malleability
- Hardness
- Magnetic anisotropy
- Coercivity
- Saturation magnetization
- Bose-Einstein condensation
- Fermi energy
- Debye temperature
- Phonon lifetime
- Thermal diffusivity
- Peltier effect
- Wiedemann-Franz law
- Magnetic field strength
- Magnetic dipole moment
- Electromotive force
- Avalanche breakdown
- Effective mass
- Landau levels
- Plasmon resonance
- Lattice constant
- Grain boundary
- Point defects
- Vacancy defects
- Schottky defect
- Thermal gradient
- Critical current density
- Quantum tunneling
- Meissner effect
- London penetration depth
- Density functional theory
- Exchange-correlation energy
- Electron-electron interaction
- Dielectric constant
- Spin relaxation
- Spin Hall effect
- Magnetoresistance
- Colossal magnetoresistance
- Skyrmions
- Topological superconductors
- Spin ice
- Charge density wave
- Orbital ordering
- Anomalous Hall effect
- Spin current
- Thermoelectric efficiency
- Electrical resistivity
- Paramagnetism
- Piezoelectricity
- Hall effect
- Neel temperature
- Work function
- Surface tension
- Inductance
- Absorption coefficient
- Spin-orbit coupling
- Hole mobility
- Permittivity
- Shear modulus
- Young's modulus
- Compressibility
- Plasticity
- Thermal expansion
- Phase transition
- Boiling point
- Viscosity
- Ductility
- Tensile strength
- Hysteresis
- Remanence
- Superfluidity
- Quantum Hall effect
- Fermi level
- Phonon dispersion
- Photon absorption
- Seebeck coefficient
- Thomson effect
- Magnetic flux density
- Spin density
- Magnetic monopole
- Zener breakdown
- Carrier concentration
- Cyclotron frequency
- Plasma frequency
- Exciton binding energy
- Crystal symmetry
- Dislocation density
- Interstitial defects
- Frenkel defect
- Diffusion coefficient
- Heat flux
- Critical temperature
- Josephson effect
- Flux pinning
- Cooper pairs
- Local density approximation
- Phonon-electron interaction
- Polarizability
- Energy dissipation
- Spin lifetime
- Rashba effect
- Giant magnetoresistance
- Spintronics
- Topological insulators
- Quantum spin liquid
- Frustrated magnetism
- Spin density wave
- Valleytronics
- Spin polarization
- Thermopower
- Lorenz number
```

```
- Mott transition                    - Hubbard model
- Anderson localization              - Kondo effect
- Heavy fermions                     - Spin glass
- Quantum entanglement               - Decoherence
- Quantum coherence                  - Quantum dots
- Single-photon emitter              - Two-photon absorption
- Quantum well                       - Quantum wire
- Dirac fermions                     - Weyl fermions
- Band topology                      - Chern number
- Berry phase                        - Berry curvature
- Topological invariants             - Luttinger liquid
- Spin-charge separation             - Fractional quantum Hall effect
- Anyons                             - Majorana fermions
- Photonic crystals                  - Metamaterials
- Hyperbolic materials               - Negative refraction
- Zero refractive index              - Anisotropic conductivity
- Isotropic conductivity             - Heat capacity
- Thermodynamic stability            - Gibbs free energy
- Helmholtz free energy              - Entropy
- Enthalpy                           - Order parameter
- Critical exponents                 - Universality class
- Phase coexistence                  - Triple point
- Solid-liquid interface             - Liquid-vapour interface
- Wetting                            - Contact angle
- Hydrophobicity                     - Hydrophilicity
- Surface roughness                  - Surface energy
- Adhesion                           - Cohesion
- Slip boundary condition            - Stokes flow
- Reynolds number                    - Knudsen number
- Mach number                        - Shock waves
- Turbulence                         - Chaos theory
- Lyapunov exponent                  - Bifurcation
- Strange attractors                 - Fractal dimension
- Atomic structure                   - Electron configuration
- Molecular bonding                  - Covalent bonding
- Ionic bonding                      - Metallic bonding
- Van der Waals forces               - Hydrogen bonding
- Crystal lattice                    - Lattice defects
- Grain size                         - Dislocations
- Impurities                         - Doping
- Alloying                           - Phase composition
- Microstructure                     - Polymorphism
- Amorphous structure                - Crystallinity
- Chemical composition               - Elemental abundance
- Atomic radius                      - Ionic radius
- Polarizability                     - Electronegativity
- Ionization energy                  - Electron affinity
- Band structure                     - Energy levels
- Density of states                  - Fermi level
- Carrier density                    - Effective mass
- Valence electrons                  - Conductive electrons
- Phonon interactions                - Electron-phonon coupling
- Spin-orbit coupling                - Spin-spin interactions
- Magnetic domains                   - Exchange interaction
- Anisotropy energy                  - Dipole-dipole interaction
- Orbital overlap                    - Hybridization
- Molecular orbitals                 - Surface states
- Quantum confinement                - Strain
- Stress                             - External pressure
- Temperature                        - Thermal history
- Quenching                          - Annealing
```

C ROC CURVES FOR EACH KEYWORD

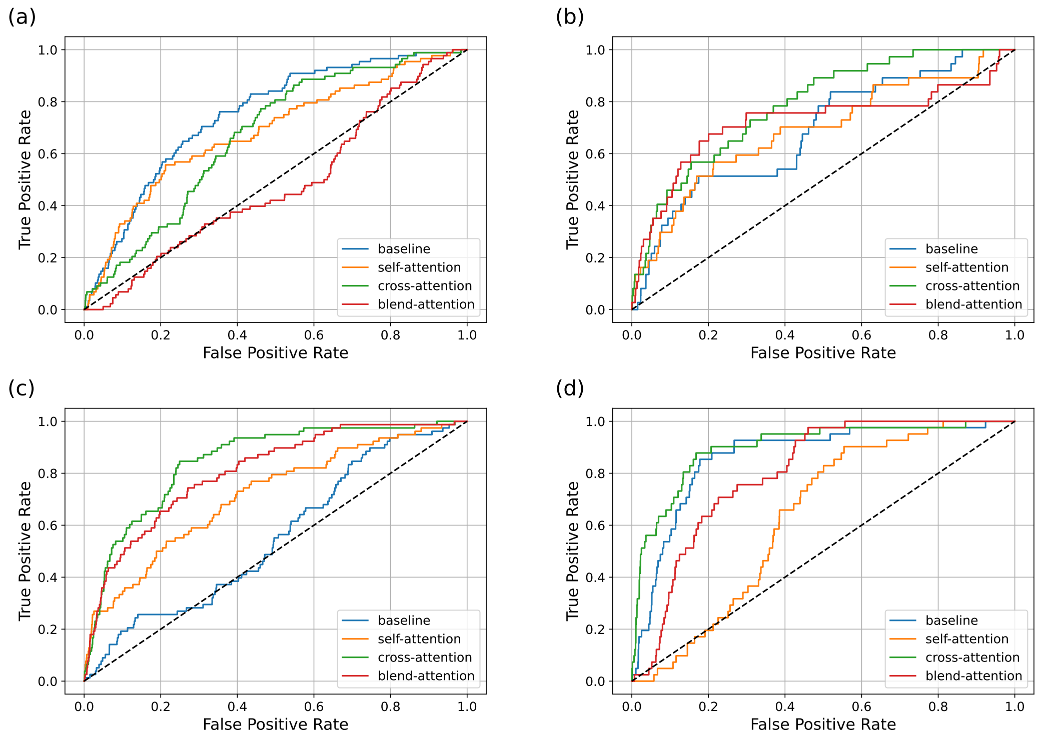

Figure A1: ROC curves of zero-shot cross-modal crystal structure classification using the keywords (a) : “ferromagnetic,” (b) : “ferroelectric,” (c) : “semiconductor,” and (d) : “electroluminescence” on the test set.