Title

Towards the holistic design of alloys with large language models

Journal

Nature Reviews Materials, 9(12)

ISSN

2058-8437

Authors

Pei, Zongrui

Yin, Junqi

Neugebauer, Jörg

et al.

Publication Date

2024-12-01

DOI

10.1038/s41578-024-00726-6

Towards the holistic design of alloys with large language models

Zongrui Pei,[1] Junqi Yin[2], Jörg Neugebauer[3], Anubhav Jain[4]

[1]New York University, New York, NY10012, USA
[2]Oak Ridge National Laboratory, Oak Ridge, TN37831, USA
[3]Max-Planck-Institut für Eisenforschung, Düsseldorf, 40237, Germany
[4]Lawrence Berkeley National Laboratory, Berkeley, CA94720, USA
E-mail: peizongrui@gmail.com

Large language models are very effective at solving general tasks, but can also be useful in materials design, extracting and using information from the scientific literature and unstructured corpora. In the domain of alloy design and manufacturing, they can expedite the materials design process and enable the inclusion of holistic criteria.

Of the many alloys that are designed and synthesized, many never find their way into practical applications or manufacturing workflows. To bridge this gap materials design must become more holistic, considering a wide array of constraints, including processing, microstructure, sustainability, and properties. This complexity requires a more integrated approach than is typically taken in current research. A more comprehensive and sophisticated strategy is needed to address the demands of metal development and production in the context of global societal and industrial needs. A single research team usually lacks the expertise to cover all these facets when designing metallic materials. To make things worse, the avalanche of scientific publications, patents, conference abstracts and other forms of corpora is not organized in a way that facilitates holistic design. Large language models (LLMs) are very well suited to extract and process information from vast corpora. Examples include OpenAI's ChatGPT, Google's Gemini, and Meta's Llama 2 and 3.1. These models may provide valuable support in the holistic design of metallic materials.

[H1] Structure-processing-property relationship

An alloy's properties are not solely dictated by its chemical composition (1–3). Predicting mechanical properties requires knowledge of chemical compositions, synthesis methods, and processing history, as processing routes and the resulting microstructures significantly alter properties. Mechanical properties and compositions can be extracted from texts using named entity recognition (NER) techniques and structured, but not all materials data can be automatically extracted and structured in this manner. This is a problem for analytical methods and machine-learning models that take only structured data. However, LLMs can incorporate such unstructured information, making them well-suited for predicting alloy properties (4). Recently, an LLM called AlloyBERT was developed to predict yield stress by incorporating text-formulated relationships, demonstrating the potential of directly using structure-processing information. Language models were also employed to design high-entropy and thermoelectric alloys (5; 6). These word-embedding studies employed the word vectors extracted from language models to find the correlations between materials and their potential properties that determine their applications. Depending on the holistic criteria, these methods can be extended into various workflows that consider multiple screening steps, each corresponding to one criterion. Although preliminary, these results showcase the potential of language models in the design of more sustainable, high-performance alloys.

[H1] Holistic alloy design with LLMs

[H2] Formulating the holistic design criteria for LLMs

One critical question is how to formulate holistic design criteria in LLMs to extract the information from publications. Here, we discuss a few possible solutions by taking sustainability as an example, suggesting prompting strategies that incorporate sustainability considerations along with requirements for specific alloy design tasks. We can use a general-purpose LLM or fine-tune a foundation model for sustainability based on publications on sustainability in metallurgy. The model can shortlist materials that have the desired properties based on their cosine similarity with keywords like "sustainability", "recycling", and "CO_2 footprint". A multi-step strategy can be developed when multiple keywords are used. For example, the above screening can be furthered based on "harmful" (to discard materials harmful to health) and "corrosion-control" (to mitigate corrosion). An alternative method is using LLMs to identify a list of environmentally friendly chemical elements. In the next step, the design space or the candidate chemical species are limited to these elements, similar to what was done for HEAs (6). Another pathway is to use LLMs to parse external information from recent journal articles or patents (rather than generating information from LLMs) and then to train other models, such as convolutional neural networks, on those data sets to generate suggestions for new materials that combine the desired alloy properties and sustainability requirements.

[H2] Processing the relationships of holistic constraints

Once we identify the holistic design criteria or constraints, the next step is to process their relationships. The holistic design of alloys is mathematically an optimization problem with multiple objectives. Given the complexity of this problem, a structured workflow is essential to achieve the goals incrementally. At each step, one objective is realized, requiring us to prioritize the constraints based on the holistic criteria of the design problem. This process involves operations research and the removal of potential conflicts among constraints. In the latter case, some objectives may be discarded. For example, many pairs of mechanical properties, such as strength and elongation to failure, exhibit trade-offs. Developing an effective strategy to manage these contradictory constraints is crucial. LLMs can help prioritize constraints, suggest potential trade-offs, and provide insights into optimal decision-making strategies, ensuring a more efficient alloy design process.

[H2] Engineering the prompts for holistic constraints

LLMs can potentially address highly complex and societally relevant issues in alloy design, but this requires well-engineered prompts. Effective prompting is crucial to obtain structured and accurate content from LLMs for metal research, as well as to avoid hallucinations. As alloys become increasingly complex in terms of chemical compositions and physical properties, the need for precise prompts grows. This involves a deep understanding of prompting theory, well-designed ontologies and hierarchies, and the identification of appropriate prompting approaches tailored to metal physics. One promising solution is to use the Retrieval-Augmented Generation (RAG) method to enhance prompt quality. RAG can mitigate hallucinations by cross-referencing outputs with quality-controlled sources like Wikipedia or scientific publications. This process ensures that the information generated by LLMs is accurate and reliable, thereby enhancing the credibility of the model's outputs.

[H2] The synergy between additive manufacturing and LLMs

Additive manufacturing (AM) is particularly promising for metallic alloys because it enables the production of complex shapes that are impossible to obtain with traditional methods. The integration of LLMs into AM can revolutionize the industry and facilitate the realization of automated holistic alloy design. LLMs can design superior metallic materials specifically suitable for manufacturing by AM and optimize the planning and strategies involved in AM products. For example, they can generate predefined instructions to control the AM process and streamline the manufacturing process.

[H1] Challenges of LLMs for metallic materials

[H2] Ambiguous alloy naming and limited standardization

Alloys can appear under multiple names in the corpora due to the reordering of chemical elements, use of hyphens, or brand and domain-specific names. This issue is especially prominent with high-entropy alloys; for example, CoCrFeMnNi can have 120 different names just by permuting the order of elements. Additionally, abbreviations are not unique; for example, ”AM” can stand for ”advanced materials” or ”additive manufacturing,” depending on the context. This inconsistency undermines the advantage of LLMs and knowledge-graph models, which rely on accurate information retrieval. Ambiguous names can prevent accurate retrieval of whether an alloy has been designed and synthesized, wasting time and resources. Standardizing names for alloys based on their elemental composition is feasible, but it becomes challenging when brand or domain-specific names are used. Therefore, there is an urgent need for a comprehensive database that includes alloy names and their brand or domain-specific equivalents. Building such a database has proven successful in other research areas in which artificial intelligence is gaining importance, such as ImageNet for deep learning and the Protein Data Bank database of protein structures for AlphaFold.

[H2] Lack of failed alloy design examples

Researchers prefer publishing positive results. High-performance alloys often drive publications, but failures in alloy design are equally important. We need to encourage researchers to report failures alongside successes rather than presenting a misleadingly smooth path to discovery. Moreover, the lack of negative results in published research and training data can bias LLMs. Fortunately, questions and feedback from metallurgists and materials scientists regarding various alloy-design failures are available in platforms like ResearchGate. Incorporating feedback to learn from failures is essential to ensure more balanced and accurate models. The community could benefit from dedicated platforms or journals for reporting failures. Structured data sets generated from automated laboratories, if shared publicly, can be another solution.

[H2] Updating LLMs

Building LLMs from scratch is costly and resource-intensive. For example, training the Llama2-70B model required 1.7 million GPU hours and produced 291 tons of CO_2 equivalent (7). Thus, it is more cost-effective to continuously update existing LLMs and fine-tune general-purpose models for metallic alloys research than to develop domain-specific models. RAG allows LLMs to access up-to-date, peer-reviewed, and quality-controlled information without needing full retraining. By combining pretrained LLMs with retrieval mechanisms, RAG enhances the quality of generated outputs for metallic alloys research, improving both material design and manufacturing processes. This approach provides a practical and efficient way to keep LLMs current and relevant in this rapidly evolving field.

[H1] Outlook

LLMs can assist in holistic material and manufacturing design tasks at various levels of involvement, ranging from straightforward data and paper summarization to agent-, simulation-, and retrieval-augmented autonomous insight and workflow development (8). In the automated design and manufacturing process, LLMs receive the instructions and transform them into operations (experiments or simulations) by selecting suitable equipment. Because each piece of equipment is an independent unit, this coordination can be challenging. When multiple targets and criteria are involved in the design of an alloy, input from different experimental and simulation tools needs to be combined. The ultimate goal for the future is to automate the whole process, from hypothesis formulation to alloy design and eventually to alloy synthesis, without any human involvement and with a fine-tuned LLM acting as a control center.

References

[1] Liu, X., Zhang, J. & Pei, Z. Machine learning for high-entropy alloys: Progress, challenges and opportunities. *Progress in Materials Science* 101018 (2022).

[2] Pei, Z. *et al.* Machine-learning microstructure for inverse material design. *Advanced Science* **8**, 2101207 (2021).

[3] Pei, Z. An overview of modeling the stacking faults in lightweight and high-entropy alloys: Theory and application. *Materials Science and Engineering: A* **737**, 132–150 (2018).

[4] Olivetti, E. A. *et al.* Data-driven materials research enabled by natural language processing and information extraction. *Applied Physics Reviews* **7**, 041317 (2020).

[5] Tshitoyan, V. *et al.* Unsupervised word embeddings capture latent knowledge from materi-

als science literature. *Nature* **571**, 95–98 (2019). URL https://doi.org/10.1038/s41586-019-1335-8.

[6] Pei, Z., Yin, J., Liaw, P. K. & Raabe, D. Toward the design of ultrahigh-entropy alloys via mining six million texts. *Nature Communications* **14**, 54 (2023).

[7] Touvron, H. *et al.* Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).

[8] Bauer, S. *et al.* Roadmap on data-centric materials science. *Modelling and Simulation in Materials Science and Engineering* (2024).

Competing interests

The authors declare no competing interests.

Acknowledgement

The authors appreciate the fruitful discussion and comments from Prof. Dierk Raabe.