Reaction Graph Networks for Inorganic Synthesis Condition Prediction of Solid State Materials

Thorben Prein[1,2,4,5], **Fuzhan Rahmanian**[1], **Kesava Prasad Arul**[1,2], **Jasmin El-Wafi**[1,2], **Menelaos Panagiotis Fotiadis**[1,2], **Jan Heimann**[1,2], **Paul Weinmann**[1,2], **Yifei Duan**[3], **Elton Pan**[3], **Elsa Olivetti**[3], **Jennifer L.M. Rupp**[1,3,5]

[1] Technische Universität München, [2] TUM.AI, [3] Massachusetts Institute of Technology, [4] Munich Data Science Institute, [5] TUMint. Energy Research GmbH
Corresponding author: jrupp@tum.de

Abstract

The integration of advanced machine learning (ML) techniques with density functional theory (DFT) has significantly enhanced the optimization and prediction of stable material structures. However, translating these computational predictions into successful laboratory syntheses—whether by autonomous labs or human scientists remains time- and cost-intensive due to the complex optimization of solid-state reaction parameters. In this work, we present the first application of a Reaction Graph Network (RGN) to model precursor interactions in inorganic reactions and predict synthesis conditions based on solid state reactions. Our approach enables the efficient prediction of synthesis conditions and demonstrates improvements over previous methods. This streamlines the path from computational predictions to material synthesis and offers potential to accelerate materials discovery.

1 Introduction

Chemical reactions forming solid state materials require the precise optimization of synthesis parameters such as temperatures, times, and environmental conditions to obtain targeted phases at sufficient yield. Among various synthesis options, solid-state reactions are fundamental to creating high-value functional materials used in energy conversion and storage, such as batteries [18, 26], as well as in opto-ionic, dielectric and magnetic applications [7, 6]. These synthesis are labor-, time- and cost-intensive and necessitate a deep understanding of the underlying reaction mechanisms and reaction parameters [15, 14]. A typical solid-state synthesis reaction of metal oxides occurs at high temperatures via solid state grain boundary and volume diffusion forming a novel compound. For the synthesis of new compounds, a deep understanding of the underlying reaction mechanisms is required [4]. Experimentally, numerous temperature settings are often tested, driving up both cost and time [32, 36, 37]. Historically, researchers have relied on heuristic methods and expert intuition to navigate these complexities [16].

While computational predictions leveraging density functional theory (DFT) and machine learning (ML) techniques [23, 24, 19] for suggesting (novel) stable structures [28, 20] have shown promise, a significant gap remains in translating these predictions into successful laboratory synthesis [1]. To bridge this gap, there is increasing interest in using high-throughput synthesis and autonomous laboratories to target these computationally derived structures within a ML-guided optimization [32, 33, 31, 22, 25]. Previously studies have explored the use of ML for prediction of synthesis conditions. Herein, Huo et al. [11] applied models such as linear regression and decision trees to predict optimal synthesis temperatures, using curated features such as precursor properties and reaction thermodynamics. Their research highlighted the predominant influence of temperature in determining

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

the success of the synthesis process. Insufficiently low or high temperatures can result in incomplete reactions, impurities, stoichiometry alterations of the chemistry or the absence of the desired target phase. Karpovich et al. [13] explored the use of a conditional variational autoencoder (CVAE) and tree-based models to classify reactions and predict reaction conditions. They find that point-regression models outperform generative models, particularly when utilizing compositional features alongside precursor melting points and anion fractions. Interestingly, Pan et al. [25] demonstrated that adopting a generative diffusion approach to materials synthesis outperforms regression and other deep generative models, as diffusion models are able to capture the one-to-many structure-synthesis relationship in zeolitic materials. While these works made significant progress, two main questions are yet to be tackled: first, how to effectively model precursor interactions within predicted reactions; and second, how to augment the small reaction datasets through the incorporation of prior chemical knowledge, e.g. in the form of pretrained models [29, 5].

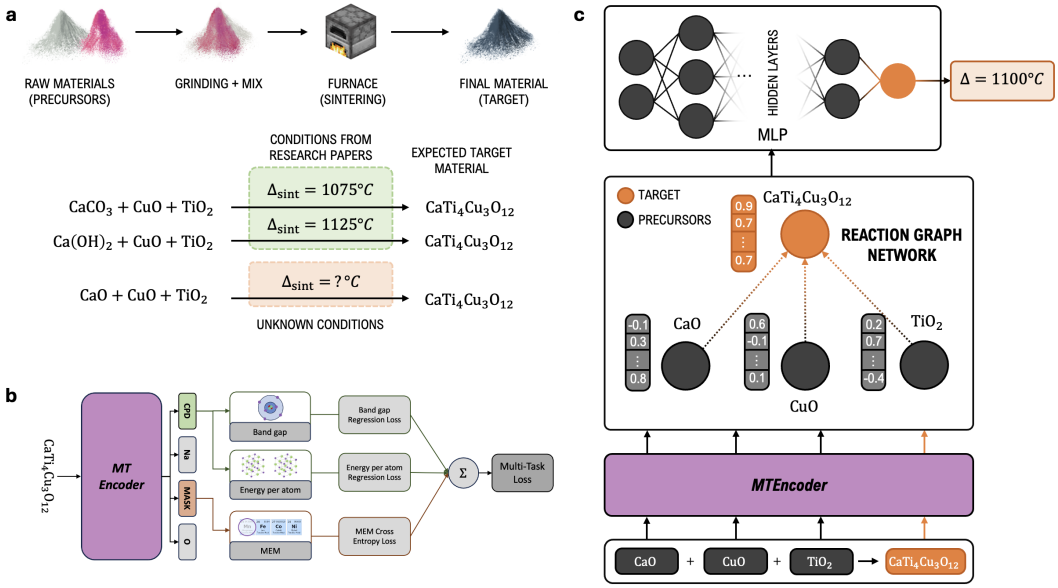

Figure 1: An overview of the integrated workflow for the proposed graph-based approach to predict (unknown) synthesis reaction conditions. a) In the classic solid-state synthesis metal oxides and/or hydroxides are reacted as educts towards a new product compound. In such a synthesis process high temperature solidification via sintering is involved. The sintering conditions, such as temperature (Δ_{sint}), are traditionally derived from literature or determined through empirical optimization. Our framework aims to guide experimental synthesis by predicting sintering temperatures. b) The proposed method begins by embedding precursor and target compounds with the structure-agnostic representations of the multi-task trained MTEncoder [27] model. c) The RGN model leverages these learned representations within a directed graph to model inorganic reactions. Reaction representations are then derived through an additive pooling operation over the graph, followed by regression using a MLP to predict reaction temperatures.

In particular, our work: 1) introduces the first Reaction Graph Neural Network (RGN) model designed to represent inorganic reactions and achieve accurate predictions of synthesis conditions and 2) extends prior research by exploring the use of deep-learned material representations for predicting reaction conditions. This has the potential to streamline the path from predicted structures to synthesized materials.

2 Methods

2.1 Structure Agnostic Representations

In synthesis-related tasks, structural information is often lacking, primarily due to how datasets are extracted from academic literature. A significant portion of these datasets features incomplete structural data, either due to non-reporting within the original publications or the inherent difficulties

associated with extracting detailed structural information from textual sources. We search for structural information regarding the compounds present in the text-mined synthesis condition dataset [17] and find only 14% of these to have a match in the Materials Project [12]. This limitation prevents the use of structural information, as the mapping between composition and crystal structure is non-injective. Meaning, a single composition can correspond to multiple stable phases, making it impossible to map directly from composition to a unique structure.

Our approach, instead, relies on structure-agnostic representations from MTEncoder (Figure 1b)[27]. This model is pretrained on 12 distinct tasks representing DFT-computed properties across 1.5 million compounds in the Alexandria database [8]. In prior work, we demonstrated that these representations generalize well to downstream tasks and achieve state-of-the-art results for structure-agnostic property prediction. For our experiments, we leverage two model variants: one trained on 512-dimensional representations and the other on 256-dimensional representations.

2.2 Reaction Graphs

We model reactions as a graph (Figure 1c), with each precursor represented by a node connected to the target compound involved in the reaction. By excluding byproducts, we ensure the applicability of our framework in a materials discovery setting and account for possible inconsistencies in the dataset, which might derive from incomplete reporting and extraction. We explore different RGN architectures and find a graph directed from precursor to target compounds to perform best (see ablation study A.4). We hypothesize that this emphasizes the causal relationships in the reaction process, ensuring the model to learns how precursor interactions and reaction conditions collectively influence the formation of the final product. Furthermore, we use graph attention network (GAT) layers to contextualize the precursor and target compounds [34], achieving improved performance. These layers aggregate property- and structure-related information from all compounds involved in the reaction. To stabilize the training process and improve gradient flow, we apply a batch normalization layer after the graph attention layers [3]. RGN employs additive pooling to obtain graph-level representations, which are then passed through a series of feedforward layers to regress the reaction temperature (Figure 3). Given the variation in graph sizes across reactions, selecting an appropriate pooling method is essential. Our ablation study indicates that additive pooling yields the best performance, suggesting it may better preserve information about the number of different precursors involved. Thus, we employ additive pooling in our approach (see A.4).

3 Experimental Setup

3.1 Datasets

Our work is based on the largest text-mined dataset of inorganic reactions available [17]. This dataset is derived by classifying synthesis procedures and conditions with language models. We process the dataset extensively to yield 8,049 entries with complete precursor and target compound data (more details in A.2). For model training and evaluation, we use a 5-fold cross-validation split of 70% training, 10% validation, and 20% testing, maximizing data utility in this relatively small dataset.

3.2 Material Embeddings

We compare our approach to three other materials embeddings. The compositional representations sparsely represent each elements fraction in a respective column. For the matminer representations [35], atomic number, Mendeleev number, atomic weight, melting point, periodic table position, covalent radius, and electronegativity are chosen in resonance with previous work [13]. The MatSciBERT model [9] is utilized to embed materials through text-based questions related to sintering temperature. To capture the compound representations we encode sentences with the respective compounds, namely *"What is the sintering temperature of [compound]?"*.

3.3 Models

For model comparison, we use two standard baselines: XGBoost and a feedforward neural network implemented in PyTorch [13]. Featurization of the reaction equations is achieved by mean pooling the precursor and target compound representations.

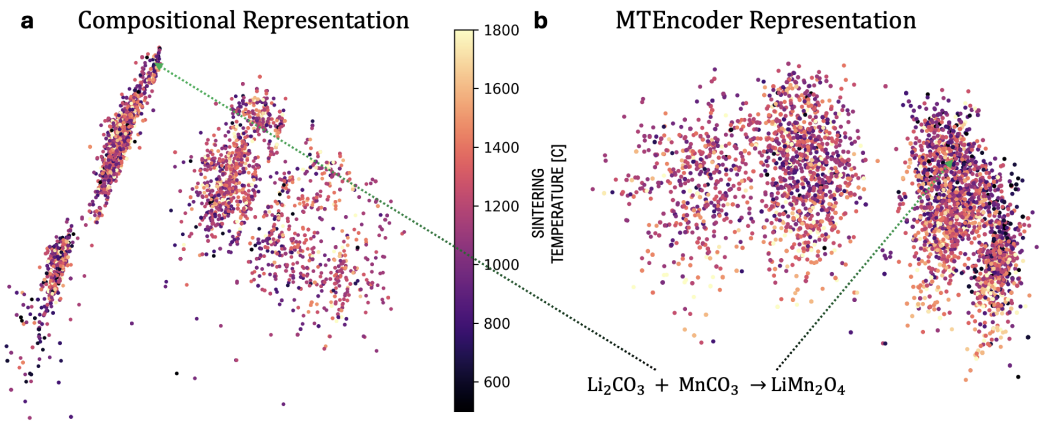

Figure 2: Comparison of reaction representations in the solid-state reaction dataset [17] using two different embedding methods. **a.** shows compositional representations, while **b.** illustrates 256-dimensional MTEncoder representations. In both cases, the precursor and target compound representations are mean pooled and then dimensionally reduced using PCA. Shown representations are not trained on the sintering temperatures. The dotted line showcases an example reaction within the dataset.

4 Results

4.1 Comparing Embeddings

The results, as shown in Table 1, indicate that the lowest MAEs are achieved using MTEncoder representations for each of the tested models. Between the 256- and 512-dimensional representations, we observe a slight improvement with the lower-dimensional representation. This may be attributed to the associated smaller number of model parameters, which can be advantageous given the limited size of the training dataset. The purely compositional representations perform second best among the different models, aligning with prior research from [13] relying on this approach. The strong performance of these representations, based solely on atomistic fractions, can be attributed to the high correlation between elemental features and synthesis processes, such as the association between high temperatures and oxides. Results for the Matminer features show an increased MAE margin of approximately 20-30 °C, likely due to the lack of compositional information about the elements in the compounds (see 3.2). Furthermore, the language model-based MatSciBERT [9] representations result in the highest MAEs. This could be due to MatSciBERT's pretraining on general materials science domains and its focus on textual meaning, which may limit its ability to capture the granularities specific to material compositions.

Overall, we observe that the embeddings learned by MTEncoder achieve the best performance across all representations, a novel finding not previously reported. The information density inherent in MTEncoder representations is reflected in Figure 2b. It is important to note that the representations shown in the figure have not yet been trained on the synthesis condition task.

4.2 Learning Compound Interactions

In addition to comparing embeddings, we evaluated the influence of model architecture. The results show that RGN models consistently outperform their NN counterparts. For example, with the MTEncoder representation, the RGN model achieved a lower MAE of 89.9 °C compared to 112 °C for the NN model. This supports our thesis that the RGN's ability to contextualize material presence and model interactions is more effective than applying mean pooling to non-contextualized representations of the compounds involved in a reaction.

Table 1: Summary of benchmarking results with different models on the sintering temperature dataset from [17], using a 5-fold split. The values represent test mean absolute errors (MAE) and their standard deviations for the reaction graph network (RGN), neural network (NN), and XGBoost (XGB) models. The best result for each model is marked in bold.

Featurization		*RGN (Ours)* (↓)	*NN* (↓)	*XGB* (↓)
Physical	Composition	98.6 (2.0)	115.3 (3.4)	159.3 (6.7)
	Matminer	120.9 (7.5)	127.4 (5.6)	173.1 (23.2)
Deep	MTEncoder (512)	90.8 (3.3)	**109.3** (4.3)	153.5 (10.1)
	MTEncoder (256)	**89.9** (4.0)	112.0 (5.2)	**152.3** (10.5)
	MatSciBERT	120.6 (3.7)	172.1 (8.7)	156.9 (22.7)

Mean baseline 221.2 (4.8) - all values in °C

4.3 Discussion

Our approach to learn sintering conditions through the RGN simplifies the synthesis problem, by assuming a one-to-one mapping between synthesis recipes and outcomes. This means having only a single valid sintering temperature for a given reaction of educts, whereas in laboratory settings, this is inherently a one-to-many relationship, with multiple temperatures potentially leading to successful synthesis. This stemms from the possibility to alter sinter times as well. Furthermore, in real-world lab settings, temperature profiles are used and linked to phase formation sequences. Unfortunately, incorporating temperature profiles and accounting for one-to-many relationships requires larger datasets for effective learning, which are currently unavailable. Another potential enhancement may be achieved by focusing on reactions with fewer reactants, thus avoiding the complexity of learning a single reaction condition for intermediate products and byproducts formed in multi-step reactions. This aligns with a recent approach for automated discovery of novel compounds by reacting two phases at a time [32]. Although data on pairwise reactions remains limited, our approach is readily transferable to this new setting.

A promising direction for model improvement could involve fine-tuning the MTEncoder model during training of the reaction graph network. A further improvement to the current approach could involve incorporating precursor quantity information, such as stoichiometric coefficients in reaction equations, which may offer valuable insights. This data could be represented, for instance, as edge attributes in the RGN, enhancing the model's capacity to capture reaction specifics. However, this would require access to stoichiometric data, which is often missing in the current dataset. Furthermore, future work could expand on the current approach by predicting additional synthesis parameters in solid-state synthesis and addressing more complex direct wet-to-solid densification methods, including condensation, sol-gel processes, and pyrolysis reactions [38, 37, 36, 21, 10, 2, 30].

5 Conclusion

We propose Reaction Graph Networks, a graph-based approach for predicting sintering conditions, benchmarking the model on a comprehensive dataset of inorganic solid-state reactions, which includes diverse materials such as functional metal oxides, used for battery cathodes. We show that 1) introducing a graph architecture to represent compound interactions and 2) the usage of deep learned features significantly improves results in inorganic synthesis condition prediction by approximately 20 °C. Additionally, our approach enables improved usability, as no prior knowledge of precursor or target properties is required and compositional information is sufficient as input. We hope this will accelerate the experimental discovery of new materials in the future.

6 Acknowledgements

The authors JLMR and TP wish to express their sincere gratitude to the Bavarian Ministry of Economic Affairs, Regional Development and Energy and TUMint.Energy Research GmbH for their generous support.

References

[1] Evan R Antoniuk, Gowoon Cheon, George Wang, Daniel Bernstein, William Cai, and Evan J Reed. Predicting the synthesizability of crystalline inorganic materials from the data of known material compositions. *npj Computational Materials*, 9(1):155, 2023.

[2] Moran Balaish and Jennifer LM Rupp. Widening the range of trackable environmental and health pollutants for li-garnet-based sensors. *Advanced Materials*, 33(20):2100314, 2021.

[3] Nils Bjorck, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. Understanding batch normalization. *Advances in neural information processing systems*, 31, 2018.

[4] Juan R Chamorro and Tyrel M McQueen. Progress toward solid state synthesis by design. *Accounts of chemical research*, 51(11):2918–2925, 2018.

[5] Pierre-Paul De Breuck, Geoffroy Hautier, and Gian-Marco Rignanese. Materials property prediction for limited datasets enabled by feature selection and joint learning with modnet. npj computational materials, 7 (1): 1–8, 2021. *Javed, A., Rehman, MA, Hassan, F., Shahzad, W., Ahmed, S., & Singh (2022). Thar Baker, and Thippa Reddy Gadekallu. Integration of blockchain technology and federated learning in vehicular (iot) networks: A comprehensive survey. Sensors (Basel, Switzerland)*, 22(12):4394, 2021.

[6] Thomas Defferriere, Ahmed Sami Helal, Ju Li, Jennifer LM Rupp, and Harry L Tuller. Ionic conduction-based polycrystalline oxide gamma ray detection–radiation-ionic effects. *Advanced Materials*, page 2309253, 2024.

[7] Thomas Defferriere, Dino Klotz, Juan Carlos Gonzalez-Rosillo, Jennifer LM Rupp, and Harry L Tuller. Photo-enhanced ionic conductivity across grain boundaries in polycrystalline ceramics. *Nature Materials*, 21(4):438–444, 2022.

[8] Mohammad M Ghahremanpour, Paul J Van Maaren, and David Van Der Spoel. The alexandria library, a quantum-chemical database of molecular properties for force field development. *Scientific data*, 5(1):1–10, 2018.

[9] Tanishq Gupta, Mohd Zaki, NM Anoop Krishnan, and Mausam. Matscibert: A materials domain language model for text mining and information extraction. *npj Computational Materials*, 8(1):102, 2022.

[10] Zachary D Hood, Yuntong Zhu, Lincoln J Miara, Won Seok Chang, Philipp Simons, and Jennifer LM Rupp. A sinter-free future for solid-state battery designs. *Energy & Environmental Science*, 15(7):2927–2936, 2022.

[11] Haoyan Huo, Christopher J Bartel, Tanjin He, Amalie Trewartha, Alexander Dunn, Bin Ouyang, Anubhav Jain, and Gerbrand Ceder. Machine-learning rationalization and prediction of solid-state synthesis conditions. *Chemistry of Materials*, 34(16):7323–7336, 2022.

[12] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.

[13] Christopher Karpovich, Elton Pan, Zach Jensen, and Elsa Olivetti. Interpretable machine learning enabled inorganic reaction classification and synthesis condition prediction. *Chemistry of Materials*, 35(3):1062–1079, 2023.

[14] Kun Joong Kim, Moran Balaish, Masaki Wadaguchi, Lingping Kong, and Jennifer LM Rupp. Solid-state li–metal batteries: challenges and horizons of oxide and sulfide solid electrolytes and their interfaces. *Advanced Energy Materials*, 11(1):2002689, 2021.

[15] Kun Joong Kim and Jennifer LM Rupp. All ceramic cathode composite design and manufacturing towards low interfacial resistance for garnet-based solid-state lithium batteries. *Energy & Environmental Science*, 13(12):4930–4945, 2020.

[16] Holger Kohlmann. Looking into the black box of solid-state synthesis. *European Journal of Inorganic Chemistry*, 2019(39-40):4174–4180, 2019.

[17] Olga Kononova, Haoyan Huo, Tanjin He, Ziqin Rong, Tiago Botari, Wenhao Sun, Vahe Tshitoyan, and Gerbrand Ceder. Text-mined dataset of inorganic materials synthesis recipes. *Scientific data*, 6(1):203, 2019.

[18] Rubayyat Mahbub, Kevin Huang, Zach Jensen, Zachary D Hood, Jennifer LM Rupp, and Elsa A Olivetti. Text mining for processing conditions of solid-state battery electrolytes. *Electrochemistry Communications*, 121:106860, 2020.

[19] Shreshth A Malik, Rhys EA Goodall, and Alpha A Lee. Predicting the outcomes of material syntheses with deep learning. *Chemistry of Materials*, 33(2):616–624, 2021.

[20] Amil Merchant, Simon Batzner, Samuel S Schoenholz, Muratahan Aykol, Gowoon Cheon, and Ekin Dogus Cubuk. Scaling deep learning for materials discovery. *Nature*, 624(7990):80–85, 2023.

[21] Willis O'Leary, Livia Giordano, and Jennifer LM Rupp. Tuning reduction conditions to understand and control ni exsolution from sr 0.8 la 0.1 ca 0.1 ti 0.94 ni 0.06 o 3- δ. *Journal of Materials Chemistry A*, 11(39):21429–21442, 2023.

[22] Willis O'Leary, Manuel Grumet, Waldemar Kaiser, Tomáš Bučko, Jennifer LM Rupp, and David A Egger. Rapid characterization of point defects in solid-state ion conductors using raman spectroscopy, machine-learning force fields, and atomic raman tensors. *Journal of the American Chemical Society*, 2024.

[23] Elton Pan, Christopher Karpovich, and Elsa Olivetti. Deep reinforcement learning for inverse inorganic materials design. *arXiv preprint arXiv:2210.11931*, 2022.

[24] Elton Pan, Soonhyoung Kwon, Zach Jensen, Mingrou Xie, Rafael Gómez-Bombarelli, Manuel Moliner, Yuriy Román-Leshkov, and Elsa Olivetti. Zeosyn: A comprehensive zeolite synthesis dataset enabling machine-learning rationalization of hydrothermal parameters. *ACS Central Science*, 10(3):729–743, 2024.

[25] Elton Pan, Kwon Soonhyoung, Sulin Liu, Mingrou Xie, Yifei Duan, Thorben Prein, Killian Sheriff, Yuriy Román, Manuel Moliner, Rafael Gómez-Bombarelli, and Elsa Olivetti. A chemically-guided generative diffusion model for materials synthesis planning. In *AI for Accelerated Materials Design - NeurIPS 2024*, 2024.

[26] Reto Pfenninger, Michal Struzik, Iñigo Garbayo, Evelyn Stilp, and Jennifer LM Rupp. A low ride on processing temperature for fast lithium conduction in garnet solid-state battery films. *Nature Energy*, 4(6):475–483, 2019.

[27] Thorben Prein, Elton Pan, Tom Doerr, Elsa Olivetti, and Jennifer LM Rupp. Mtencoder: A multi-task pretrained transformer encoder for materials representation learning. In *AI for Accelerated Materials Design-NeurIPS 2023 Workshop*, 2023.

[28] Patrick Reiser, Marlen Neubert, André Eberhard, Luca Torresi, Chen Zhou, Chen Shao, Houssam Metni, Clint van Hoesel, Henrik Schopmans, Timo Sommer, et al. Graph neural networks for materials science and chemistry. *Communications Materials*, 3(1):93, 2022.

[29] Nima Shoghi, Adeesh Kolluru, John R Kitchin, Zachary W Ulissi, C Lawrence Zitnick, and Brandon M Wood. From molecules to materials: Pre-training large generalizable models for atomic property prediction. *arXiv preprint arXiv:2310.16802*, 2023.

[30] Philipp Simons, Steven A Schenk, Marco A Gysel, Lorenz F Olbrich, and Jennifer LM Rupp. A ceramic-electrolyte glucose fuel cell for implantable electronics. *Advanced Materials*, 34(24):2109075, 2022.

[31] Helge S Stein and John M Gregoire. Progress and prospects for accelerating materials science with automated and autonomous workflows. *Chemical science*, 10(42):9640–9649, 2019.

[32] Nathan J Szymanski, Pragnay Nevatia, Christopher J Bartel, Yan Zeng, and Gerbrand Ceder. Autonomous and dynamic precursor selection for solid-state materials synthesis. *Nature communications*, 14(1):6956, 2023.

[33] Nathan J Szymanski, Yan Zeng, Haoyan Huo, Christopher J Bartel, Haegyeom Kim, and Gerbrand Ceder. Toward autonomous design and synthesis of novel inorganic materials. *Materials horizons*, 8(8):2169–2198, 2021.

[34] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

[35] Logan Ward, Alexander Dunn, Alireza Faghaninia, Nils ER Zimmermann, Saurabh Bajaj, Qi Wang, Joseph Montoya, Jiming Chen, Kyle Bystrom, Maxwell Dylla, et al. Matminer: An open source toolkit for materials data mining. *Computational Materials Science*, 152:60–69, 2018.

[36] Yuntong Zhu, Michael Chon, Carl V Thompson, and Jennifer LM Rupp. Time-temperature-transformation (ttt) diagram of battery-grade li-garnet electrolytes for low-temperature sustainable synthesis. *Angewandte Chemie*, 135(45):e202304581, 2023.

[37] Yuntong Zhu, Zachary D Hood, Haemin Paik, Pedro B Groszewicz, Steffen P Emge, Farheen N Sayed, Chengjun Sun, Moran Balaish, David Ehre, Lincoln J Miara, et al. Highly disordered amorphous li-battery electrolytes. *Matter*, 7(2):500–522, 2024.

[38] Yuntong Zhu, Ellis R Kennedy, Bengisu Yasar, Haemin Paik, Yaqian Zhang, Zachary D Hood, Mary Scott, and Jennifer LM Rupp. Uncovering the network modifier for highly disordered amorphous li-garnet glass-ceramics. *Advanced Materials*, 36(16):2302438, 2024.