

MIT Open Access Articles

Inorganic Materials Synthesis Planning with Literature-Trained Neural Networks

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation:

Published Version: 10.1021/ACS.JCIM.9B00995

Publisher: American Chemical Society (ACS)

Permanent Link: <https://hdl.handle.net/1721.1/132534>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: <http://creativecommons.org/licenses/by-nc-sa/4.0/>



Inorganic Materials Synthesis Planning with Literature-Trained Neural Networks

Edward Kim,¹ Zach Jensen,¹ Alexander van Grootel,¹ Kevin Huang,¹ Matthew Staib,² Sheshera Mysore,³ Haw-Shiuan Chang,³ Emma Strubell,³ Andrew McCallum,³ Stefanie Jegelka,² and Elsa Olivetti^{1,*}

¹Dept. of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

²Dept. of EECS and CSAIL, Massachusetts Institute of Technology, Cambridge, MA, USA

³College of Information and Computer Sciences,
University of Massachusetts Amherst, Amherst, MA, USA

(Dated: February 17, 2019)

Leveraging new data sources is a key step in accelerating the pace of materials design and discovery. To complement the strides in synthesis planning driven by historical, experimental, and computed data, we present an automated method for connecting scientific literature to synthesis insights. Starting from natural language text, we apply word embeddings from language models, which are fed into a named entity recognition model, upon which a conditional variational autoencoder is trained to generate syntheses for arbitrary materials. We show the potential of this technique by predicting precursors for two perovskite materials, using only training data published over a decade prior to their first reported syntheses. We demonstrate that the model learns representations of materials corresponding to synthesis-related properties, and that the model’s behavior complements existing thermodynamic knowledge. Finally, we apply the model to perform synthesizability screening for proposed novel perovskite compounds.

Recent advances in predicting material properties [1–3], screening synthesizable compounds [4–6], and organic reaction prediction [7–9] have been driven, in part, by the accessibility of machine-readable datasets [10–12] and, consequently, data-driven models. In stark contrast to organic reaction databases [12], the overwhelming majority of *inorganic* synthesis knowledge lies locked within the text of journal articles [13] and laboratory notebooks [14]. While the latter has been shown as an effective source for guiding successful syntheses, there is currently no general framework for automatically drawing synthesis insights from the literature at large.

Although scientific literature has previously been used to illuminate patterns in nanoscale morphologies [13], device performances [15], and apparatus parameters [16], each of these efforts have required tailored data representations and algorithms. To rapidly translate literature knowledge into synthesis planning resources without fine-tuned adjustments, transfer learning across different materials systems is necessary [13]. Thus, leveraging data from broad volumes of scientific literature is a critical step towards capturing synthesis trends and extending them to unknown materials.

In this work, we present an automated method for connecting scientific literature to insights for materials synthesis planning. We show that an unsupervised conditional variational autoencoder (CVAE) [17–19] can generate synthesis predictions for a variety of materials, including materials unknown to the model. This CVAE learns directly from the materials synthesis literature and produces an internal representation of precursors which corresponds to physical and chemical trends without receiving any explicit domain knowledge. We use the liter-

ature knowledge captured by the CVAE to complement first-principles techniques in materials screening tasks.

To accelerate the efforts of the materials science community, we open-source several key resources used in this work [20]: We release context-sensitive embeddings from language models (ELMo) that have been adapted for materials science text [21] along with a pre-trained FastText word embedding model for materials science [22]. Each of these embedding models has been trained on a collection of over 2.5 million materials science journal articles [13, 23]. Finally, we provide over two hundred annotated literature synthesis routes for named entity recognition (NER) tasks, such as identifying reaction conditions and materials.

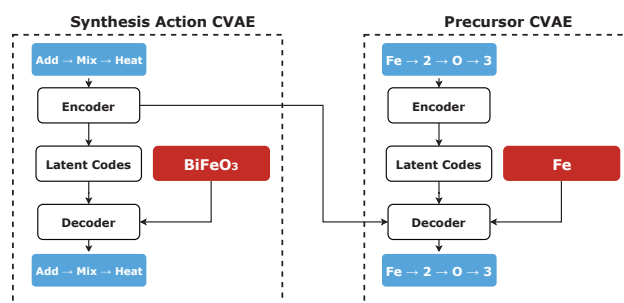


FIG. 1. Schematic diagram of the CVAE architecture. The model consists of two joined CVAEs used for learning synthesis actions and precursors, respectively. The Synthesis Action CVAE learns distributions of synthesis actions sequences conditioned on target materials. The Precursor CVAE learns distributions of precursor formulas conditioned on both a target element and an encoded representation of the jointly-observed synthesis action sequence. Target materials are represented by FastText embeddings, and all other inputs to the model are sequences of one-hot vectors.

* Corresponding author: elsao@mit.edu

We first describe our automated workflow. After a recurrent neural network [24] identifies synthesis sections of journal articles, context-sensitive ELMo word embeddings are computed and passed into another recurrent neural network which performs NER to identify precursors, synthesis target materials, and synthesis actions. Then, a CVAE model, shown in Figure 1, is trained to learn representations of synthesis routes from the named entities in an unsupervised manner. More details on these methods are provided in the Supplementary Methods.

To maximize the opportunity for transfer learning of synthesis trends, we choose a broad definition for synthesis routes that requires minimal assumptions. For a given target material m , a synthesis route S_m^i is a 2-tuple consisting of a sequence of n synthesis actions (a_1, a_2, \dots, a_n) acting on a set of l precursors $\{p_1, \dots, p_l\}$,

$$S_m^i = ((a_k)^n, \{p_j\}^l) \quad (1)$$

and in general, a single target material m may have $N > 1$ valid synthesis routes and thus S_m^i represents the i th valid synthesis route for m . We also define precursors p as “element sources,” such that they are materials sharing an element with m . The CVAE model is then constructed to model the following distributions,

$$\mathbb{P}((a_k)^n | \theta_a, m) \quad (2)$$

$$\mathbb{P}(p_j | \theta_p, e_j, (a_k)^n) \quad (3)$$

where θ_a and θ_p are model parameters for the synthesis action and precursor CVAEs, respectively, and e_j is the shared element between a precursor p_j and target material m (e.g., titanium).

Since CVAEs are generative models, novel synthesis actions and precursors can be generated by sampling from a Gaussian prior distribution [17]. Critically, we represent m by FastText word embeddings which enable the transfer of synthesis trends between existing and novel materials by leveraging literature-based similarity.

TABLE I. Generated precursors for InWO_3 and PbMoO_3 , drawn from the CVAE model. The CVAE model was trained on synthesis routes published during or before 2005. A more detailed version is available in the Supplemental Results.

| Target Material | Precursors |
|------------------|---|
| InWO_3 | $\text{In}_2\text{S}_3 + \text{WCl}_4$ |
| | $\text{In}(\text{NO}_3)_3 + \text{WCl}_4$ |
| | $\text{In}_2\text{O}_3 + \text{WO}_2$ |
| | $\text{In}_2\text{O}_3 + \text{WN}$ |
| | $^\dagger \text{InCl}_3 + \text{Na}_2\text{WO}_4$ |
| PbMoO_3 | $\text{PbCl}_2 + \text{MoCl}_2$ |
| | $\text{PbSO}_4 + \text{MoCl}_2$ |
| | $^\ddagger \text{PbO} + \text{MoO}_2$ |


[†] Precursors match Kamalakkannan et al. (2016) [25].

[‡] Precursors match Takatsu et al. (2017) [26].


To demonstrate the applicability of our CVAE method, we construct a dataset of approximately 51,000 synthe-

sis action sequences and 116,000 precursors via a general set of search terms (“perovskite + thermoelectric + multiferroic + photovoltaic + solar + nano + cathode”) and apply our neural network pipeline. We investigate the effectiveness of the CVAE model in synthesis planning by performing a publication-year-split experiment, where the model is trained only on syntheses published prior to 2005 (~ 2800 syntheses). We apply the model in predicting precursors for materials that were unseen during training, are computationally predicted as stable perovskites [1], and only recently appear in the literature: InWO_3 and PbMoO_3 , first reported in 2016 and 2017, respectively [25, 26]. Table I shows a report of the data generated by sampling from the CVAE’s Gaussian prior distributions, where the CVAE suggests the precursors for both materials (see Table ?? for additional details). The CVAE model is thus capable of predicting synthesis precursors while relying only on literature knowledge from more than a decade prior to the literature-reported syntheses of these materials. Trial-and-error (or random) precursor selection is substantially less efficient, as the number of possible precursor sets for each material is in the hundreds. Thus literature-driven models may greatly accelerate future synthesis attempts of novel materials.

During the data generation process, the CVAE model proposes several plausible syntheses beyond the literature-matching samples. To the best of the authors’ knowledge, the only reported synthesis of InWO_3 is via a solution-phase route. However, the CVAE model suggests that solid state synthesis of InWO_3 may be possible, using In_2O_3 and either WO_2 or WN as precursors. Such syntheses may be feasible, as they are thermodynamically favorable (using data at 0 K and 0 atm from OQMD) [10]:



$$\Delta H = -158 \text{ kJ/mol}$$



$$\Delta H = -930 \text{ kJ/mol}$$

We do note, however, that these simple thermodynamic analyses can only be used as rough guidelines. Besides the limitations of estimating an overall thermochemical reaction for the synthesis, along with extrapolation from STP conditions, kinetic effects are not considered here. Indeed, while it is common to mix binary oxide precursors in solid state syntheses of ternary (or quaternary, etc.) oxides, the use of nitride precursors is far less common due to the high bond energies of many nitride compounds [27]. To achieve a clearer understanding of kinetic effects, experimental verification would be required alongside a model which incorporates reaction conditions (e.g., temperatures), and this is an area for future work.

In the suggested recipes for PbMoO_3 , the CVAE model suggests a solution-phase route using PbSO_4 and MoCl_2 , both of which are soluble under acidic conditions. This may provide another viable path towards synthesizing

PbMoO₃, which has only been realized so far by solid state synthesis methods [26]. Thus, the CVAE may be used as a source of suggestions for synthesis planning, although we stress that these suggestions still need human evaluation, further analysis, and cannot be applied “out of the box.”

Despite the fact that chemical knowledge is never given to the CVAE model, solubility rules emerge from the model results. To demonstrate this, we generate W-bearing precursors for InWO₃ conditioned on two representative action sequences sampled from the Synthesis Action CVAE: A solid-state synthesis (mix, grind, calcine, press, sinter, cool) and a solution-phase synthesis (add, dissolve, stir, heat, wash, dry). Following this, we generate 10,000 CVAE-suggested precursors. The most common W-bearing precursor generated for the solid-state synthesis is the water-insoluble WO₃, while the most common precursor generated for the solution-phase synthesis is the highly-soluble Na₂WO₄. The differences of these precursor likelihoods in each case is substantial, with -16% and +21% changes to the likelihoods of the CVAE suggesting WO₃ and Na₂WO₄, respectively, when switching from conditioning on solid-state to solution-phase synthesis actions.



This effect of learning precursor trends from the literature is further demonstrated upon inspecting latent codes learned by the model. Since the CVAE learns conditional distributions, the input precursors are projected into a degenerate latent space, where the degeneracy is split by the conditional input received by the decoder. By investigating several examples (see Figure ??), we find that the CVAE learns to group precursors with similar synthesis-relevant properties, including insoluble binary oxides, water-soluble polyanion compounds, and pure/alloyed metals. This suggests that the CVAE model is capable of capturing chemical intuition and composition-driven similarity solely by joint observations of precursors, synthesis actions, and target materials. Despite the lack of “negative” data in the literature, the diversity of published synthesis literature is sufficient to drive the CVAE model in learning meaningful representations of precursors.

To emphasize the particular nature of synthesis planning via a *literature-trained* model, we contrast suggested precursors by the CVAE model with thermodynamic stability computations from OQMD [10]. Figure 2a shows a graph representation of the chemical space spanned by all CVAE-suggested precursors, where graph vertices are precursors and graph edges represent thermodynamic two-phase equilibria [28]. Clearly, the CVAE is not simply suggesting all thermodynamically-viable precursors, as there is a substantial set of precursors never suggested by the CVAE (blue vertices versus red vertices). Indeed, the CVAE has filtered precursors from a set of 73 possible precursors to only 27. Additionally, Figures 2b and 2c show that the CVAE is not selecting precursors in correspondence with isolated thermodynamic metrics: the CVAE’s suggestions are explained neither by thermody-

amic reactivity (i.e., the number of relevant two-phase equilibria with respect to other precursors) nor individual precursor stability (i.e., formation energy). This suggests that there is a meaningful difference between the thermodynamically-driven and literature-driven synthesis planning methods. While the former probes the realm of physical possibility, the latter emphasizes practical choices and historical trends.

We next train the CVAE model on our full dataset, using no publication-year cutoffs. To investigate the capability of the model for suggesting syntheses of a novel, never-before-synthesized material, we consider novel ABO₃ perovskite materials proposed by Balachandran et al. [29]. These proposed perovskites have not previously been synthesized, and have high thermodynamic stability as measured by energy differences against their convex hulls. We note that ABO₃ perovskites are used here as a representative example due to their chemical variety and diverse range of properties, but the CVAE model does indeed generalize to other categories of materials (see Table ??).

HgZrO₃ is one such example of a thermodynamically stable, unsynthesized perovskite material [29], and we perform synthesis predictions using the CVAE model (see Table ??). We find that the CVAE proposes solid-state syntheses which appear to be thermodynamically reasonable:



Again, we emphasize that thermodynamic analyses are often insufficient to evaluate reaction plausibility. As an additional utility for evaluating generated synthesis parameters, we develop a similarity metric based on the latent codes learned by the CVAE. By measuring nearest-neighbors of latent codes for the recipe using mercuric oxide and zirconium carbide, we find that the two closest literature recipes are for solid-state syntheses of SrZrO₃ and BaAl₂O₄ (see Figure ??). Besides providing insight into which observed literature examples “inspired” this particular prediction, we are also led to further insights on precursor selections. ZrC is an uncommon choice of precursor, but carbonate precursors are readily used in solid-state syntheses. Indeed, both of the near-neighbor syntheses for HgZrO₃ use carbonate precursors rather than carbides.

While similarity methods have previously been produced for materials (e.g., based on crystal structures) [30], the CVAE incorporates synthesis knowledge to produce a distinct measure of similarity. Indeed, from a structural point of view, it would not be expected that SrZrO₃ and BaAl₂O₄ should have high similarity to HgZrO₃, since all three materials form ground-

dwel times. This additional domain knowledge may be incorporated by filtering proposed synthesis parameters [7] or constraining model outputs [31]. Motivated by these possibilities, our open-source NER annotations include the necessary labels (e.g., reaction conditions) to enable these future studies.

We would like to acknowledge funding from the National Science Foundation Award 1534340, DMREF that provided support to make this work possible, support

from the Office of Naval Research (ONR) under Contract No. N00014-16-1-2432, the MIT Energy Initiative, and NSF CAREER #1553284. Early work was collaborative under the Dept. of Energy’s Basic Energy Science Program through the Materials Project under Grant No. EDCBEE. This work was also partly funded by the MIT-Sensetime Alliance on Artificial Intelligence. We would also like to acknowledge valuable feedback from Rafael Jaramillo, Gerbrand Ceder, Olga Kononova, Wenhao Sun, Haoyan Huo, and Tanjin He.

-
- [1] T. Xie and J. C. Grossman, *Physical Review Letters* **120**, 145301 (2018).
- [2] T. D. Huan, A. Mannodi-Kanakkithodi, and R. Ramprasad, *Physical Review B* **92**, 014106 (2015).
- [3] O. Isayev, D. Fourches, E. N. Muratov, C. Oses, K. Rasch, A. Tropsha, and S. Curtarolo, *Chemistry of Materials* **27**, 735 (2015).
- [4] B. Meredig, A. Agrawal, S. Kirklin, J. E. Saal, J. Doak, A. Thompson, K. Zhang, A. Choudhary, and C. Wolverton, *Physical Review B* **89**, 094104 (2014).
- [5] M. Aykol, V. I. Hegde, S. Suram, L. Hung, P. Herring, C. Wolverton, and J. S. Hummelshøj, arXiv preprint arXiv:1806.05772 (2018).
- [6] K. Kim, L. Ward, J. He, A. Krishna, A. Agrawal, and C. Wolverton, *Physical Review Materials* **2**, 123801 (2018).
- [7] M. H. Segler, M. Preuss, and M. P. Waller, *Nature* **555**, 604 (2018).
- [8] H. Gao, T. J. Struble, C. W. Coley, Y. Wang, W. H. Green, and K. F. Jensen, *ACS Central Science* (2018).
- [9] C. W. Coley, R. Barzilay, T. S. Jaakkola, W. H. Green, and K. F. Jensen, *ACS Central Science* **3**, 434 (2017).
- [10] J. E. Saal, S. Kirklin, M. Aykol, B. Meredig, and C. Wolverton, *Jom* **65**, 1501 (2013).
- [11] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, *et al.*, *APL Materials* **1**, 011002 (2013).
- [12] J. Goodman, “Computer software review: Reaxys,” (2009).
- [13] E. Kim, K. Huang, A. Saunders, A. McCallum, G. Ceder, and E. Olivetti, *Chemistry of Materials* **29**, 9436 (2017).
- [14] P. Raccuglia, K. C. Elbert, P. D. Adler, C. Falk, M. B. Wenny, A. Mollo, M. Zeller, S. A. Friedler, J. Schrier, and A. J. Norquist, *Nature* **533**, 73 (2016).
- [15] L. Ghadbeigi, J. K. Harada, B. R. Lettiere, and T. D. Sparks, *Energy & Environmental Science* **8**, 1640 (2015).
- [16] S. R. Young, A. Maksov, M. Ziatdinov, Y. Cao, M. Burch, J. Balachandran, L. Li, S. Somnath, R. M. Patton, S. V. Kalinin, *et al.*, *Journal of Applied Physics* **123**, 115303 (2018).
- [17] D. P. Kingma and M. Welling, in *International Conference on Learning Representations* (2014).
- [18] K. Sohn, H. Lee, and X. Yan, in *Advances in Neural Information Processing Systems* (2015) pp. 3483–3491.
- [19] R. Gómez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik, *ACS Central Science* **4**, 268 (2018).
- [20] www.github.com/olivettigroup/materials-synthesis-generative-models.
- [21] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)* (2018).
- [22] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, *Transactions of the Association for Computational Linguistics* **5**, 135 (2017).
- [23] E. Kim, K. Huang, A. Tomala, S. Matthews, E. Strubell, A. Saunders, A. McCallum, and E. Olivetti, *Scientific Data* **4**, 170127 (2017).
- [24] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, in *NIPS 2014 Workshop on Deep Learning, December 2014* (2014).
- [25] J. Kamalakkannan, V. Chandraboss, and S. Senthilvelan, *World Scientific News* **58**, 97 (2016).
- [26] H. Takatsu, O. Hernandez, W. Yoshimune, C. Prestipino, T. Yamamoto, C. Tassel, Y. Kobayashi, D. Batuk, Y. Shibata, A. M. Abakumov, *et al.*, *Physical Review B* **95**, 155105 (2017).
- [27] W. Sun, A. Holder, B. Orvañanos, E. Arca, A. Zakutayev, S. Lany, and G. Ceder, *Chemistry of Materials* **29**, 6936 (2017).
- [28] V. I. Hegde, M. Aykol, S. Kirklin, and C. Wolverton, arXiv preprint arXiv:1808.10869 (2018).
- [29] P. V. Balachandran, A. A. Emery, J. E. Gubernatis, T. Lookman, C. Wolverton, and A. Zunger, *Physical Review Materials* **2**, 043802 (2018).
- [30] L. Yang and G. Ceder, *Physical Review B* **88**, 224107 (2013).
- [31] M. J. Kusner, B. Paige, and J. M. Hernández-Lobato, in *International Conference on Machine Learning* (2017) pp. 1945–1954.
- [32] F. Chollet *et al.*, “Keras,” (2015).
- [33] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, *et al.*, in *OSDI*, Vol. 16 (2016) pp. 265–283.
- [34] S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson, and G. Ceder, *Computational Materials Science* **68**, 314 (2013).
- [35] M. C. Swain and J. M. Cole, *Journal of chemical information and modeling* **56**, 1894 (2016).
- [36] S. Katyayan and S. Agrawal, *Journal of Materials Science: Materials in Electronics* **28**, 18442 (2017).
- [37] A.-K. Larsson, R. Withers, J. Perez-Mato, J. F. Gerald, P. J. Saines, B. J. Kennedy, and Y. Liu, *Journal of Solid State Chemistry* **181**, 1816 (2008).