

Optimizing Fractional Compositions To Achieve Extraordinary Properties

Andrew R. Falkowski, Steven K. Kauwe, and Taylor D. Sparks

May 24, 2021

Traditional, data-driven materials discovery involves screening chemical systems with machine learning algorithms and selecting candidates that excel in a target property. The number of screening candidates grows infinitely large as the fractional resolution of compositions the number of included elements increases. The computational infeasibility and probability of overlooking a successful candidate grow likewise. Our approach shifts the optimization focus from model parameters to the fractions of each element in a composition. Using a pretrained network, CrabNet, and writing a custom loss function to govern a vector of element fractions, compositions can be optimized such that a predicted property is maximized or minimized. Single and multi-property optimization examples are presented that highlight the capabilities and robustness of this approach to inverse design.

1 Introduction

Humanity faces an ever-present need for the development of new materials with enhanced properties to tackle the energy, climate, communication, health, and security challenges of the future. However, serendipity has been the rule, not the exception, when it comes to materials discovery [1]. It is critical that materials scientists develop new ways to rationally explore chemical white space in order to more efficiently and systematically unlock new materials

The Materials Genome Initiative (MGI) was proposed for this exact purpose ten years ago. The concept behind this initiative was that big data and data science could potentially accelerate materials discovery by reducing the number of costly simulations and experiments [2]. In some ways the MGI has proven wildly successful! Successful case studies abound showcasing the power of high fidelity predictive machine learning (ML) models working in concert with modeling and experiment. For example, the discovery of the first 3D printable aluminum alloy [3], superhard materials [4], and record setting perovskite solar cells [5]. The successes of these examples are due in large part to the expansive, and growing databases for materials properties [6–11] but also to advanced ML tools ranging from feature generation (CGCNN [12], jarvis [13], oliynyk, magpie) to algorithms (RooSt [14], CrabNet [15]).

Even as materials informatics becomes increasingly mainstream and common as a research tool, limitations and criticisms of the technique mount. For example, despite theoretical evidence that extrapolation to exceptional materials is possible [16], using ML to identify truly new and unexpected materials has been relatively elusive. Another criticism is that materials informatics is truly a remarkable tool for rapid materials *screening* but a relatively poor tool for *discovery* since many of the best ML approaches rely on rapid prediction of properties from a large list of *known* materials as opposed to a generative ML model that would generate new *unknown* suggestions for compounds and compositions of interest. Another problem is that models do a great job of predicting a materials property for a given composition or compound, but at present are incapable of doing the reverse; suggesting an exact composition or compound that would have a target property. Instead, researchers settle for screening from either known materials or a large number of handpicked candidate compositions. This approach is fatally flawed when one considers that chemical white space is on the order of 10^{12} even when we restrict considerations to only quaternary systems composed of stable elements with modest granularity (>0.03).

Inverse design aims to tackle this problem by starting with targeted properties and working backwards to find a suitable chemical formula. Designing systems that are capable of understanding and mapping the the vast possible chemical space is an ongoing challenge. The hope is that a system can leverage learned representations of materials to generate novel compositions with extraordinary properties.

Neural style transfer has been demonstrated as a computationally inexpensive means of generating data from learned representations. The method, first introduced by Gatys et al. [17] learns a style representation of an artistic work and a content representation for a photograph by extracting information from a pretrained convolutional neural network at specific layers. The style and content representations for a white noise image are extracted at the same layers. A loss function for each representation is crafted such that the difference between the style and content representations of the white noise and those of the painting and picture, respectively, are minimized. By computing the associated gradient, the pixel values of the white noise image can be updated such that the style and content loss functions are minimized. The result is an image with the content of the photograph in the style of the painting. An example of the Mona Lisa (content) merged with the The Great Wave off Kanagawa (style) is shown in Figure 1.

This approach leverages a model’s learned understanding of data to generate useful data. With the availability of highly accurate ML models in the materials informatics space, it is likely that this method could be adapted to generate compositions with targeted properties. This approach deviates from present implementations as the inputs to a model become the parameters that are to be optimized.

In this work the model parameters of a trained network are frozen and the focus of optimization is shifted to the fractional components of compositions. A custom loss function is constructed such that element fractions can be updated through gradient computation to generate compositions with maximized or minimized predicted properties.



Figure 1: An image with the content of the Mona Lisa by Leonardo da Vinci and the style of The Great Wave off Kanagawa by Katsushika Hokusai produced with neural style transfer by P.J. Finlay.

2 Methods

Adapting style transfer to materials informatics follows a similar but simplified approach with regards to the loss function. A pretrained network was selected that provides accurate predictions across many different properties. For this implementation, the CrabNet architecture from Wang et al. was chosen. This model provides state of the art performance in predicting from elemental information and includes self-attention mechanisms. The attention mechanisms are of interests as they better capture relationships between elemental species and provide a degree of interpretability.

Within the CrabNet featurization scheme, compositions are initially represented as an element and fraction vector. The element vector contains the atomic numbers of the chemical species; and the fraction vector contains the respective fractions of each element within the compound. These vectors are encoded such that they take on useful representations for machine learning. The element vectors are featurized according to each element’s properties using traditional *mat2vec*, *oliynyk*, *jarvis*, or *magpie* embeddings. The fractions undergo an alternating sine and cosine expansion as described by Vaswani et al. [18] and implemented by Wang et al. [15] in CrabNet. It is important to note that log operations are applied to half of the fractional encodings to better preserve the influence of dopants on the predicted properties. In the initial version of CrabNet, these fractional encodings were accessed through a table lookup. Fractional optimization requires that the gradient is preserved throughout the prediction process, so the model had to be modified such that the expanded representation is calculated on request. These elemental and fractional representations are finally merged to form an elemental descriptor matrix (EDM) that is passed to the network.

CrabNet can be divided into two main components. The first is a self-attention block that updates the elemental representations in the EDM layer based on element’s relationships to one another and to the target property. After several iterations of updates, the new representation is passed through a fully connected residual neural network. A more detailed description and discussion of CrabNet the architecture can be found in the paper from Wang et al. [15].

The initial representations of compositions in the CrabNet scheme are the focal point for adapting the style transfer approach. Rather than updating network parameters, as is the case during model training, the network weights are frozen and a new set of tunable parameters in the form of the fraction vector are chosen. The fraction vector is then analogous to the white noise image in traditional style transfer and is similarly updated by computing the gradient of a loss function. Rather than extracting style information at different layers, the loss function in this implementation seeks to maximize or minimize the final property value by minimizing the difference between the predicted value and positive or negative infinity respectfully. In the code implementation, a value of 100,000




Figure 2: Architecture for performing fractional optimization on an elemental system. Fraction vector values are treated as tunable parameters that are updated such that a predicted property is maximized or minimized. Parameter updates are governed by computing the gradient through standard error back-propagation.

is used as an analogue for infinity (see [Github](#) for more details on implementation). A diagram of the architecture is shown in Figure 2.

In this initial implementation a simple $L1$ loss function governs the updating fractional amounts. The tunable fraction vector is represented as \vec{f} . To ensure that the fractional amounts sum to one, a softmax function is applied to \vec{f} to produce an updated vector in the form of \vec{f}' . The EDM representation of \vec{f}' is then passed through the network $f()$ to produce a property prediction \hat{y} . The predicted value has been scaled to the mean of the target values in the training set. This scaling is preserved for \hat{y} as it results in better performance in multi-objective optimization. The loss function can then be written as in Eq (1).

$$\begin{aligned}
 \vec{f}' &= \text{softmax}(\vec{f}) \\
 \hat{y} &= f(\text{EDM}^{\vec{f}'}) \\
 \mathcal{L}(\vec{f})_{\text{prediction}} &= |\infty - \hat{y}|
 \end{aligned} \tag{1}$$

The gradient with respect to the fraction vector \vec{f} can thus be computed using standard error back-propagation - allowing the fraction vector to be updated such that the predicted property is maximized or minimized.

In a multi-objective optimization problem, the predictions from two trained networks are used to update the fraction vector. Two $L1$ loss functions govern each predicted property. These are summed to form a larger loss function, from which, the gradient can be calculated. Through the inclusion weighing parameters, the trade-off between two properties can be balanced according to researcher intent. This loss function is shown in Eq (2).

$$\begin{aligned}\hat{y} &= f(\text{EDM}^{\vec{f}'}) & \hat{z} &= g(\text{EDM}^{\vec{f}'}) \\ \mathcal{L}(\vec{f})_{prediction} &= \alpha \cdot |\infty - \hat{y}| + \beta \cdot |\infty - \hat{z}|\end{aligned}\quad (2)$$

3 Results and Discussion

Optimizing a chemical system for a single property provides initial validation for the developed technique. Bulk modulus is a traditionally difficult material property to predict and can serve as an analogue to material hardness as shown by Mansouri et al [4]. The Ti-Cd-C chemical system was optimized using the previously discussed method. Figure 3a shows the bulk modulus response to changing the fractional amounts of elements in the system. The increase in property response and convergence to a solution demonstrates the model's capabilities. Figure 3b shows the changes for each element in atomic percent throughout the optimization process. The fractions are initialized in a one to one ratio for all chemical species under investigation. It is noted that the algorithm attempts to remove Cd from the system such that TiC is produced - a material known for having a high bulk modulus.

The initial investigation presented in Figure 3 highlights some important aspects of the model, as it is presently implemented. Firstly, TiC was not present in the training or validation data, so the model's ability to predict this material provides indication that the model is capable of extrapolating elemental relationships and exploring chemical spaces it is unfamiliar with. Secondly, it is noted that the concentration of cadmium remains at dopant concentrations even over long optimization periods. This is a consequence of the softmax function applied to fraction vector and CrabNet's fractional embedding, which preserves the influence of dopants on a chemical system. Thus, while the fraction of Cd is small relative to the other species in the optimized composition, it continues to influence the final properties, hence why the predicted bulk modulus of the nearly TiC is lower




Figure 3: Single property optimization for the Ti-Cd-C system. (a) shows changes in predicted bulk modulus over the course of the optimization process against the calculated loss. Uncertainty in prediction is included to show range of possible values. (b) shows changes in atomic percent for each element in the system over the course of optimization.

than it's true value. It should be mentioned that when optimization is run in the absence of Cd, the prediction of TiC bulk modulus is closer to the true value.

The necessity of selecting a chemical system for optimization retains the initial problem of having to search many candidate systems. To test the robustness of the method, an additional exploration of a system with a large number of elements was explored.

Attention networks provide a level of interpretability to the model by providing information on what the algorithm pays attention to during training and the decision making process. CrabNet allows the average contribution of each element to a trained property to be quantified. Ten elements with the highest contributions to bulk modulus were selected to investigate the ability for the optimization algorithm to optimize a large chemical space. The results are shown in Figure 4.

Figure 4a shows an increasing bulk modulus response to the model's adjustments to the fractional components of each element. The percent increase in bulk modulus over the course of optimization is small relative to the results in Figure 3. Figure 4b demonstrates the model's ability to prioritize candidates among many options in order to improve performance. However, dopant influence remains and is likely heightened by the increased number of candidates.

A final validation demonstrates the ability for the model to optimize two competing properties. With only chemical information at the model's disposal, it is likely to predict compounds that are not thermodynamically stable. Bartel et al. demonstrated decom-




Figure 4: Single property optimization for a ten element system. Elements were selected based on highest average contribution to bulk modulus during model training. (a) shows changes in predicted bulk modulus over the course of the optimization process against the calculated loss. Uncertainty in prediction is included to show range of possible values. (b) shows changes in atomic percent for each element in the system over the course of optimization. The model demonstrates the capability to reduce the influence of unnecessary elements.

position energy as a reasonable analogue to thermodynamic stability and showed that networks with attention mechanisms can predict it with reasonable accuracy [19]. For this demonstration, the model is asked to maximize bulk modulus and minimize decomposition energy such that that a stable material can be found with a high bulk modulus. Figure 2 shows the results of this multi-optimization problem.

The plots in Figure 2 show more extreme behavior compared to previous examples. The predicted bulk modulus changes drastically, while the decomposition energy remains relatively constant after an initial decrease. The uncertainty in predicted properties also varies more significantly over the course of optimization. These results indicate the the method is capable of finding a balance between properties. That said, it appears that bulk modulus is prioritized during optimization - indicating an imbalance in the loss functions. While the predicted values are normalized to their respective training datasets, bulk modulus can more readily meet the loss target compared to decomposition energy.

The multi-optimization results also highlight the model's ability to correct optimization mistakes and avoid local minima. Figure 2b shows the fractional amount of calcium initially increasing before being reduced. The capacity for the algorithm to 'course correct' highlights its robustness to the common trappings in optimization problems.

The deficiencies in this present implementation of the method provide many opportunities for future work. Adjusting the architecture such that the influence of dopant concentra-




Figure 5: Plots highlighting algorithms capabilities for balancing multiple properties during optimization. (a) shows changes in predicted bulk modulus over the course of the optimization process against the calculated loss. Uncertainty in prediction is included to show range of possible values. (b) shows changes in atomic percent of each elemental species. Additionally, it shows the ability for the model to course correct so as to better optimize properties.

tions can be better controlled could provide better performance when screening a large number of candidate elements. The preservation of dopant influence by the CrabNet architecture provides advantages and disadvantages with regards to navigating a general composition space. Implementing a fractional optimization method only requires that the gradient be preserved throughout the prediction process - allowing the method to be adapted to a wide array of already trained models. Selecting models are the well suited to a property or composition space could yield better results and better capture the nuances of that space. The results presented in this paper were produced with a 3.1 GHz Dual-Core Intel Core i5 in under a minute, making the method computationally inexpensive and highly accessible.

Improvements in performance could also come from modifications to the loss function. Finding prediction balance in multi-objective optimization tasks is a common challenge. The present implementation uses a simple linear combination of weighted losses to provide the user with a degree of control over what the model prioritizes. Implementing a more sophisticated loss function could yield optimizations that correspond user expectations.

4 Conclusions and Future Directions

The results presented in this paper represent an initial investigation into a new approach to inverse design. By shifting the focus of optimization from model parameters to the fractional components of compositions, it is demonstrated that composition spaces can be navigated more efficiently to find materials with desirable properties. While the present implementation has some deficiencies in terms of general performance, further development and application specific tailoring could yield significant improvements. As fractional optimization is adaptable to many existing models, it may serve as an evaluation tool in assessing how models represent and evaluate complex chemical systems. To the authors' knowledge, this approach is yet unseen in the materials informatics space. With time, it is the hope of the authors that this approach can be further developed into a useful tool for inverse design.

5 Code Availability

The code for executing fraction optimization and producing the figures shown in this manuscript is available on GitHub here: github.com/AndrewFalkowski/CoCoCrab

References

- (1) Graser, J.; Kauwe, S. K.; Sparks, T. D. Machine learning and energy minimization approaches for crystal structure predictions: A review and new horizons. *Chemistry of Materials* **2018**, *30*, 3601–3612.
- (2) De Pablo, J. J.; Jones, B.; Kovacs, C. L.; Ozolins, V.; Ramirez, A. P. The materials genome initiative, the interplay of experiment, theory and computation. *Current Opinion in Solid State and Materials Science* **2014**, *18*, 99–117.
- (3) Martin, J. H.; Yahata, B. D.; Hundley, J. M.; Mayer, J. A.; Schaedler, T. A.; Pollock, T. M. 3D printing of high-strength aluminium alloys. *Nature* **2017**, *549*, 365–369.
- (4) Mansouri Tehrani, A.; Oliynyk, A. O.; Parry, M.; Rizvi, Z.; Couper, S.; Lin, F.; Miyagi, L.; Sparks, T. D.; Brzoch, J. Machine learning directed search for ultraincompressible, superhard materials. *Journal of the American Chemical Society* **2018**, *140*, 9844–9853.
- (5) Lu, S.; Zhou, Q.; Ouyang, Y.; Guo, Y.; Li, Q.; Wang, J. Accelerated discovery of stable lead-free hybrid organic-inorganic perovskites via machine learning. *Nature communications* **2018**, *9*, 1–8.

(6) Ghadbeigi, L.; Harada, J. K.; Lettieri, B. R.; Sparks, T. D. Performance and resource considerations of Li-ion battery electrode materials. *Energy & Environmental Science* **2015**, *8*, 1640–1650.

(7) Jain, A.; Ong, S. P.; Hautier, G.; Chen, W.; Richards, W. D.; Dacek, S.; Cholia, S.; Gunter, D.; Skinner, D.; Ceder, G.; Persson, K. a. The Materials Project: A materials genome approach to accelerating materials innovation. *APL Materials* **2013**, *1*, 011002.

(8) Calderon, C. E.; Plata, J. J.; Toher, C.; Oses, C.; Levy, O.; Fornari, M.; Natan, A.; Mehl, M. J.; Hart, G.; Nardelli, M. B., et al. The AFLOW standard for high-throughput materials science calculations. *Computational Materials Science* **2015**, *108*, 233–238.

(9) Saal, J. E.; Kirklin, S.; Aykol, M.; Meredig, B.; Wolverton, C. Materials design and discovery with high-throughput density functional theory: the open quantum materials database (OQMD). *Jom* **2013**, *65*, 1501–1509.

(10) Gaulois, M. W.; Sparks, T. D.; Borg, C. K.; Seshadri, R.; Bonificio, W. D.; Clarke, D. R. Data-driven review of thermoelectric materials: performance and resource considerations. *Chemistry of Materials* **2013**, *25*, 2911–2920.

(11) Tehrani, A. M.; Ghadbeigi, L.; Brgoch, J.; Sparks, T. D. Balancing mechanical properties and sustainability in the search for superhard materials. *Integrating Materials and Manufacturing Innovation* **2017**, *6*, 1–8.

(12) Xie, T.; Grossman, J. C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical review letters* **2018**, *120*, 145301.

(13) Choudhary, K.; Garrity, K. F.; Reid, A. C.; DeCost, B.; Biacchi, A. J.; Walker, A. R. H.; Trautt, Z.; Hattrick-Simpers, J.; Kusne, A. G.; Centrone, A., et al. The joint automated repository for various integrated simulations (JARVIS) for data-driven materials design. *npj Computational Materials* **2020**, *6*, 1–13.

(14) Goodall, R. E.; Lee, A. A. Predicting materials properties without crystal structure: Deep representation learning from stoichiometry. *Nature Communications* **2020**, *11*, 1–9.

(15) Wang, A. Y.-T.; Kauwe, S. K.; Murdock, R. J.; Sparks, T. D. Compositionally-Restricted Attention-Based Network for Materials Property Prediction: CrabNet, ChemRxiv, 2020.

(16) Kauwe, S. K.; Graser, J.; Murdock, R.; Sparks, T. D. Can machine learning find extraordinary materials? *Computational Materials Science* **2020**, *174*, 109498.

(17) Gatys, L. A.; Ecker, A. S.; Bethge, M. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp 2414–2423.

(18) Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; Polosukhin, I. Attention is all you need. *arXiv preprint arXiv:1706.03762* **2017**.

(19) Bartel, C. J.; Trewartha, A.; Wang, Q.; Dunn, A.; Jain, A.; Ceder, G. A critical examination of compound stability predictions from machine-learned formation energies. *npj Computational Materials* **2020**, *6*, 1–11.