High-Dimensional Discrete Bayesian Optimization with Self-Supervised Representation Learning for Data-Efficient Materials Exploration

Masaki Adachi
University of Oxford
Toyota Motor Corporation
masaki@robots.ox.ac.uk

Abstract

A material exploration model based on high-dimensional discrete Bayesian optimization is introduced. Features were extracted from a large-scale database of ab-initio calculations by self-supervised representation learning. Material exploration was carried out based on 100 prior target values from 6,218 candidate materials. As a baseline, ten human experts of materials science were selected and evaluated their exploration efficiency. Under the same conditions, the proposed discrete algorithm was 1.93 times as efficient as human experts on average, while the conventional continuous algorithm could not outperform them.

1 Introduction

In environmental issues, the discovery of innovative materials can bring a fundamental solution. Renewables with more efficient power conversion and storage can all be realised if better materials are found [1]. However, the exploration of materials has been like finding a needle in a haystack [2]. Humanity has discovered only a handful of practical materials so far due to the time-consuming synthesis process. Furthermore, the amount of experimental databases available is limited [3]. The necessity to balance between exploration and exploitation of experimental data inevitably leads to the setting of active learning and Bayesian optimization [4]. Improvement in the efficiency of materials exploration algorithms could accelerate the solution to our energy problems.

In contrast, theoretical calculations, known as DFT (Density Functional Theory), have made it possible to predict the material properties and explore materials in silico [5]. However, the DFT predictability is still low for some physical property values, such as a bandgap [6]. Thus, constructing the methodology to get the most out of the DFT database, which is large-scale but not accurate, is the key to enhancing materials exploration efficiency. In addition to this, there are two rooms for improvement: discrete exploration space and baseline. Previous research assume that the descriptors and performance of materials are continuous variables [7]. However, as only stable crystals can exist, both material variables should be essentially discrete [8], which could pave the way for more efficient exploration. Furthermore, most current research uses random search as a baseline [2, 4, 5, 7, 8], but a better baseline for comparison would be the performance of human experts.

In this paper, the following questions were examined:

- Can the efficiency of material exploration be improved by exploiting a large-scale DFT database with low accuracy?
- Is it more efficient for Bayesian optimization to assume discrete variables than conventional continuous variables?
- Is Bayesian optimization more efficient than human experts?

NeurIPS 2021 AI for Science Workshop.

Table 1: Algorithms of SSRL and materials exploration

Algorithm 1: SSRL	**Algorithm 2**: Materials Exploration
Input: CIFs (crystal structure) **Output**: 12-dim features	**Prior information** $\mathbf{X} \leftarrow$ 100 samples $\times$ 12-dim features $\mathbf{Y} \leftarrow$ 100 samples $\times$ 1-dim true labels $\mathbf{C} \leftarrow$ 6,118 samples $\times$ 12-dim features $\mathbf{Y}_{target} \leftarrow$ 2.534 eV
1. Load the pretrained CGCNN model 2. **While** (epoch < max_epoch) 2.1 Train(CGCNN) $\leftarrow$ 35,180 pseudo-labels 3. Replace the last layer with the identity mapping 4. **for** CIF out of 6,218 candidates 4.1 128-dim features $\leftarrow$ FeatureExtract(CIF) 4.2 12-dim features $\leftarrow$ PCA(128-dim features)	1. **While** (n_trial < max_trial) 1.1 $\bar{\mathbf{Y}} \leftarrow \boldsymbol{L}_{poisson}(\mathbf{Y}, \mathbf{Y}_{target})$ 1.2 model $\leftarrow$ ALEBO-GP($\mathbf{X}$, $\bar{\mathbf{Y}}$) 1.3 NextCandidate $\leftarrow$ ThompsonSampling($\mathbf{C}$, model) 1.4 $\mathbf{X}_{new}$, $\mathbf{Y}_{new} \leftarrow$ QueryToHiddenDatabase(NextCandidate) 1.5 append $\mathbf{X}_{new}$, $\mathbf{Y}_{new}$ to $\mathbf{X}$, $\mathbf{Y}$

The corresponding answers are as follows:

- Large-scale DFT database enhanced materials exploration efficiency by being exploited as pseudo-labels for self-supervised representation learning to extract essential features from the crystal information of materials.
- Discrete Bayesian optimization based on Thompson Sampling with a high-dimensional Gaussian process and Poisson loss was twice as efficient compared with the conventional continuous Bayesian optimization.
- In the statistical experiments with human experts in materials science, the discrete Bayesian optimization identified better materials than eight experts out of 10 within 30 trials.

The code is available on GitHub: `https://github.com/ma921/BanditMaterialsExplorer`

2 Problem Setting

A simplified problem was set to quantify the materials exploration efficiency under the same conditions for both algorithms and human experts: explore the material having as close a bandgap to 2.534 eV as possible, exploiting 100 prior target values from 6,218 candidate materials. The true labels are 6,218 experimentally-measured values, and the pseudo-labels are 35,180 direct bandgaps calculated by DFT calculations using VASP (Vienna Ab initio Simulation Package) [9].

The exploration was done within one hour, and the number of trials was 30. The number of trials was iterated when the algorithm or human queried the true value of the material's bandgap in the hidden database. Human experts could refer to general information such as lattice parameter, ionic radius, and electronegativity by searching via the supporting software (see GitHub). However, it was forbidden to search for papers or to consult with colleagues for the sake of fairness to the algorithms. The author supervised the experiments on human subjects to ensure no possible unforeseen factors, such as cheating or preparation for this test. Ten human subjects were randomly selected from the engineers and scientists working for Toyota Motor Corporation. Six were experts in semiconductors exploration, and four were experts in materials science, although not in semiconductors.

3 Discrete Bayesian Optimization

The algorithm consists of two stages, as described in Table 1. The first stage, SSRL (Self-Supervised Representation Learning) [10] was adopted to learn representations helpful for predicting the bandgap. SSRL transforms the 6,218 material candidates into 12-dimensional features. At the second stage, the materials exploration algorithm samples the next candidate material using 100 prior target values from 6,218 candidates.

Firstly, for SSRL, CGCNN (Crystal Graph Convolutional Neural Network) [11] was employed. CGCNN pioneered graph convolutional network architecture for crystals by learning representations from atomic information written in CIF (Crystallographic Information File) [12] as graph structures. In this study, a CGCNN model pretrained by the Materials Project [13] was adopted. The Materials Project is the most extensive database comprised of CIFs and DFT-calculated properties. However, the quality of the database is low, containing results mixed from altered DFT conditions and imaginary CIFs that cannot exist in the real world. As such, an attempt was made to develop a higher-quality database, by preparing 35,180 pseudo-labels of direct bandgaps for SSRL, in order to permit consistent VASP calculation. This self-made dataset was divided into a ratio of train/validation/test = 70/15/15. As a result, the prediction accuracy of the pseudo-labels for the test dataset had an MAE (Mean Absolute Error) of 0.134 eV. The trained model was then used as a feature extractor to extract the 128 dimension features in the last layer. The features were then reduced to 12 dimensions by PCA (Principal Component Analysis) [14].

Secondly, for materials exploration algorithm, the following three methods were adopted: Poisson Loss [15], ALEBO-GP (Adaptive Linear Embedding Bayesian Optimization Gaussian Process) [16], and Large-scale Thompson Sampling using CIQ (Contour Integral Quadrature) [17]. Each of these methods was employed for a different purpose.

Poisson loss was used to adapt the reward function for discrete variables, unlike conventional continuous reward functions (e.g., MAE). Poisson loss is expressed as follows:

$$\boldsymbol{L}_{poisson}(y, \hat{y}) = \frac{1}{n}\sum_{i=1}^{n}\left\{\hat{y}^i - y^i \log\left(\hat{y}^i\right)\right\} \tag{1}$$

where $\hat{y}$ is the predicted value, y is the true value and n is the number of data points.

ALEBO-GP was introduced to deal with high-dimensional Bayesian optimization. This model can efficiently explore the high-dimensional space by adaptively restricting the exploration subspace to linear space. However, higher dimensions are still challenging due to the curse of dimensionality [18]. On the other hand, the prediction of the bandgaps becomes less accurate with lower dimensions. Thus, balancing this trade-off between reducing dimensions for exploration and increasing dimensions for prediction is essential [19]. As a result of experiments, 12-dimensional features were optimal.

Thompson Sampling was introduced to allow the sampling of discrete candidate points. Furthermore, the adoption of CIQ was enabled to speed up the covariance matrix calculation, a computational bottleneck, making it possible to deal with a large-scale dataset.

The whole algorithm of discreteBO (discrete Bayesian optimization) will now be explained. Firstly, the 6,218 CIFs were transformed into a 12-dimensional features by SSRL. Next, the bandgaps of 100 prior target values were used as the true values by calculating Poisson loss. Subsequently, the ALEBO-GP were trained on 100 prior target values; the resulting model predicts the posteriors of 6,118 (6,218 - 100) candidate materials. Finally, the Thompson Sampling and CIQ selected the next candidate to minimize Poisson loss. This algorithm was repeated 30 times iteratively. An absolute error was employed as a metric of how close the selected materials were to the target value.

4 Results

The results of these experiments are shown in Figure 1. The vertical axis represents the absolute error from the target value of 2.534 eV, and the horizontal axis represents the number of trials. The black, blue, green, and red lines refer to random search results, performance of human experts, continuous Bayesian optimization [7], and discrete Bayesian optimization, respectively. The solid lines represent the mean values, and the coloured areas represent the 95% confidence intervals. Figure 1 shows that the proposed method (discreteBO) was able to find the best material on average. Data efficiency of an algorithm to be evaluated against discreteBO was evaluated by the following equation.

$$\eta_{exploration} = \frac{30}{\mathrm{argmin}\left|\mathrm{MAE}_{discreteBO} - \mathrm{MAE}^{30th}\right|} \tag{2}$$

where $\mathrm{MAE}_{discreteBO}$ is the array of MAEs of discreteBO, MAE^{30th} is the 30th MAE of an algorithm to be evaluated. As a result, the best model was 7.25, 2.42, 1.93 times as data efficient as

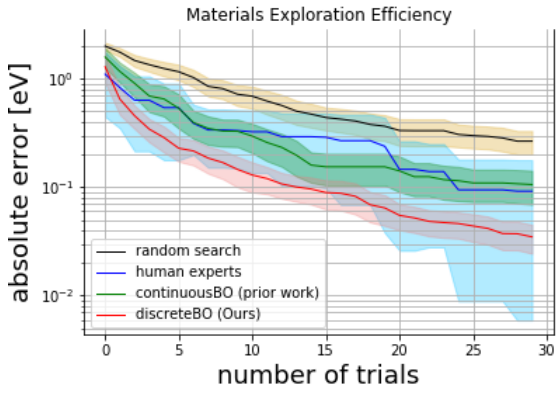

Figure 1: Materials exploration efficiency averaged over 50 times.

Table 2: Ablation Study

#	Method	MAE at 30th [eV]
1	discreteBO (Best)	**0.0348**
2	discreteBO replaced with conventional CIF discriptors from SSRL	0.1102
3	discreteBO replaced with SSRL of 0.314 MAE from 0.134 MAE	0.0780
4	discreteBO replaced with inputs of 3-dim features from 12-dim features	0.0753
5	discreteBO replaced with MAE loss from the Poisson loss	0.1070
6	discreteBO replaced with linear GP from ALEBO-GP	0.1010
7	discreteBO replaced with UCB from Thompson Sampling	0.0971
8	continuousBO	0.1060
9	Averaged results of human experts	0.0917

random search, continuousBO, and human experts, respectively. In contrast, the human variance was considerable: some individuals had the same efficiency as the random search, whereas others were more efficient than even discreteBO. Specifically, descreteBO was able to explore the dataset more efficiently than 8 out of 10 human experts. This result will require further testing in the future with an increased number of subjects.

5 Discussion

Table 2 illustrates the results of the ablation study. The difference between SSRL and conventional CIF descriptors [7] is whether or not a large-scale DFT database was used in the feature extraction stage. By comparing algorithm numbers 1 and 2, it is clearly shown that SSRL enhanced the exploration efficiency significantly (corresponding to the answer of the above question 1). Furthermore, the result of number 3 indicates that the higher accuracy of SSRL can contribute to the downstream exploration task. Moreover, number 4 explains that the 12 dimension is the optimal choice to balance between exploration and exploitation trade-off in this case. In the same way, the accuracy deterioration in the numbers 5, 6, and 7 clearly shows the proposed three components in the second learning stage are all essential, supporting the hypothesis that discrete assumptions can enhance efficiency compared to continuous ones (the above question 2).

6 Broader Impact

As there has been no baseline for materials exploration efficiency of human experts, we believe that this dataset has the potential to become a standard dataset for competing material search algorithms. Furthermore, it has also been shown experimentally that, under the right conditions, the machine learning algorithm is more efficient than human experts at discovering materials. This result could lead to an acceleration of materials discovery and potential solutions to environmental issues.

References

[1] M. Kittner, F. Lill, and D. M. Kam. Energy storage deployment and innovation for the clean energy transition. *Nature Energy*, 2(17125), 2017. URL https://www.nature.com/articles/nenergy2017125.

[2] G. Hautier. Finding the needle in the haystack: Materials discovery and design through computational ab initio high-throughput screening. *Computational Materials Science*, 163(108), 2019. URL https://www.sciencedirect.com/science/article/abs/pii/S0927025619301156.

[3] A. Zakutayev, N. Wunder, M. Schwarting, J. D. Perkins, R. White, K. Munch, W. Tumas, and C. Phillips. An open experimental database for exploring inorganic materials. *Scientific Data*, 5(180053), 2018. URL https://www.nature.com/articles/sdata201853.

[4] T. Ueno, T. D. Rhone, Z. Hou, T. Mizoguchi, and K. Tsuda. Combo: An efficient bayesian optimization library for materials science. *Materials Discovery*, 4(18), 2016. URL https://www.sciencedirect.com/science/article/abs/pii/S2352924516300035.

[5] R. Ramprasad, R. Batra, G. Pilania, A. M Kanakkithodi, and C. Kim. Machine learning in materials informatics: recent applications and prospects. *npj Computational Materials*, 3(54), 2017. URL https://www.nature.com/articles/s41524-017-0056-5.

[6] M. Yu, S. Yang, C. Wu, and N. Marom. Machine learning the hubbard u parameter in dft+u using bayesian optimization. *npj Computational Materials*, 6(180), 2020. URL https://www.nature.com/articles/s41524-020-00446-9.

[7] R. Jalem, K. Kanamori, I. Takeuchi, M. Nakayama, H. Yamasaki, and T. Saito. Bayesian-driven first-principles calculations for accelerating exploration of fast ion conductors for rechargeable battery application. *Scientific Reports*, 8(5845), 2018. URL https://www.nature.com/articles/s41598-018-23852-y.

[8] Y. Zhang, D. W. Apley, and W. Chen. Bayesian optimization for materials design with mixed quantitative and qualitative variables. *Scientific Reports*, 10(4924), 2020. URL https://www.nature.com/articles/s41598-020-60652-9.

[9] V. Eyert, O. M. Christensen, W. Wolf, D. Reith, A. Mavromaras, C. Freeman, and E. Wimmer. Unravelling the potential of density functional theory through integrated computational environments: Recent applications of the vienna ab initio simulation package in the medea® software. *Computation*, 6(63), 2018. URL https://doi.org/10.3390/computation6040063.

[10] Z. Feng, C. Xu, and D. Tao. Self-supervised representation learning by rotation feature decoupling. In *CVPR*, 2019. URL https://openaccess.thecvf.com/content_CVPR_2019/html/Feng_Self-Supervised_Representation_Learning_by_Rotation_Feature_Decoupling_CVPR_2019_paper.html.

[11] T. Xie and J. C. Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical Review Letters*, 120(45301), 2018. URL https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.120.145301.

[12] I. D. Brown and B. McMahon. The crystallographic information file (cif). *Data Science Journal*, 5(174), 2006. URL http://doi.org/10.2481/dsj.5.174.

[13] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson. The materials project: A materials genome approach to accelerating materials innovation. *APL Materials*, 1(011002), 2013. URL https://aip.scitation.org/doi/10.1063/1.4812323.

[14] K. F. R. S. Pearson. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(559), 1901. URL https://doi.org/10.1080/14786440109462720.

[15] Z. Li, J. F. Cai, and K. Wei. Toward the optimal construction of a loss function without spurious local minima for solving quadratic equations. *IEEE Transactions on Information Theory*, 66(3242), 2020. URL https://ieeexplore.ieee.org/document/8918236.

[16] B. Letham, R. Calandra, A. Rai, and E. Bakshy. Re-examining linear embeddings for high-dimensional bayesian optimization. In *NeurIPS*, 2020. URL https://arxiv.org/abs/2001.11659.

[17] G. Pleiss, M. Jankowiak, D. Eriksson, A. Damle, and J. R. Gardner. Fast matrix square roots with applications to gaussian processes and bayesian optimization. In *NeurIPS*, 2020. URL https://arxiv.org/abs/2006.11267.

[18] R. Bellman. Dynamic programming. *Science*, 153(34), 1966. URL https://www.science.org/lookup/doi/10.1126/science.153.3731.34.

[19] R. G. Bachrach, A. Navot, and N. Tishby. An information theoretic tradeoff between complexity and accuracy. *Learning Theory and Kernel Machines. Lecture Notes in Computer Science*, 2777, 2003. URL https://doi.org/10.1007/978-3-540-45167-9_43.