More trustworthy Bayesian optimization of materials properties by adding humans into the loop

Anonymous Author(s)
Affiliation
Address
email

Abstract

Bayesian optimization (BO) has proven to be effective approach for guiding sample-efficient exploration of materials domains and is increasingly being used in automated materials optimization set-ups. However, when exploring novel materials, sample quality may vary unexpectedly, which can even invalidate the optimization procedure if it remains undetected. This issue limits the use of highly-automated optimization loops, especially in high-dimensional materials spaces with a lot of samples. Sample quality may be hard to define unequivocally for a machine but human scientists are usually good at judging sample quality, at least on a cursory yet often sufficient level. In this work, we demonstrate that humans can be added into the BO loop as experts to comment on the sample quality, which results in more trustworthy BO results. We implemented human-in-the-loop BO via a data fusion approach and applied virtual BO cycles on experimental perovskite film stability data from literature. The human-in-the-loop approach facilitates automated materials design and characterization by reducing the occurrence of invalid optimization results.

1 Introduction

Bayesian optimization (BO) is an active learning method designed for globally optimizing a black-box function with sequential acquisitions of samples [1]. BO is intended especially for guiding the acquisitions, which can be costly, towards the most promising areas of the search space [3]. In particular, BO utilizes a probabilistic surrogate model (typically, a Gaussian process, GP) that predicts the target property across a fixed search space of input variables. The surrogate model is improved iteratively by collecting new samples and the sampling strategy is determined by an acquisition function. The acquisition function is based on the current surrogate model, and aims at balancing exploration of new areas or exploitation of already sampled regions [3].

BO is a promising method to be used in self-driving materials laboratories because it can be combined with automation to form semi- or fully autonomous materials optimization loops. In materials science, BO has successfully been used for guiding experimental or computational exploration of novel materials spaces (concept shown in step i in Figure 1a) and for determining optimal combinations of device settings [3-6].

One of the factors limiting the broader use of BO in experimental materials science is the varying quality of the experimentally prepared samples. The samples should be of sufficient quality to ensure valid characterization results for the target property, which is a necessity in any experimental work with or without BO. Especially when exploring materials compositions, tuning the sample preparation conditions (e.g., annealing temperature) as a function of the composition may be required for maintaining the good sample quality across the design space. Following the sample quality and adjusting procedures as necessary is possible with small sample sizes, or search spaces with existing

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

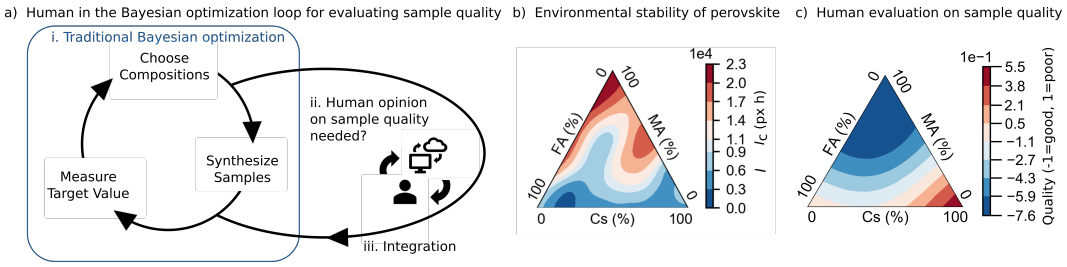

Figure 1: a) Bayesian optimization (BO) of materials samples (i) with a human added into the loop for evaluating the quality of synthesized samples. This requires a query policy (ii) and the integration (iii) of the human query result into the loop. The simulated BO loops in this work sample from two Gaussian process models (assumed as the ground truth for simulations): b) the environmental stability of perovskites (BO target) via a proxy variable I_c of sample RGB color change in pixels integrated over aging test duration in hours, and c) human evaluation on the quality of the perovskite samples fitted based on the pictures of the fresh samples. Human evaluations are shown on a scale of [-1,1]. The data for fitting the models b and c is from the literature [3].

domain knowledge on the nature and location of the possible sample quality issues. However, the increasing level of laboratory automation facilitates a shift toward higher-dimensional (i.e, more samples needed) and more novel search spaces (i.e., more unexpected events during the search). Human scientists can no longer manually check all samples in-situ, else they would severely slow down the otherwise accelerated optimization loop. As another approach, it would be principally possible to implement automated sample quality tracking via a series added characterization methods, but this is in many applications prohibitively costly and difficult to maintain. Also, sample quality is hard to define precisely and unequivocally to a machine because there is a multitude of potential sample quality issues, especially in novel search spaces in which unanticipated events frequently occur. An experienced human scientist is more versatile at detecting by quick inspection – and then confirming with a suitable characterization method when needed – sample quality issues, also the unexpected ones.

We therefore propose to add humans into the BO loop for evaluating sample quality but querying the human only when needed, thus preventing the humans from becoming a bottleneck. Human-in-the-loop approaches do exist in machine learning [7-11] but they are less frequent in materials science [12-13]. In this contribution, we show via a simulated example that is based on real experimental data [3] that adding human into the BO loop can reduce the likelihood of converging to erroneous optima with low-quality samples. Thus, adding the human into the loop increases trust in BO results with a limited amount of added effort. By showcasing the human-in-the-loop approach, we hope to trigger further discussion on what are the roles of human scientists in AI-guided design of new materials, and on how sample quality should be addressed in partially for fully self-driving materials laboratories.

2 Results and discussion

2.1 Setup

We investigate the effects of the human-in-the-loop BO in an example setting of compositional stability optimization of perovskite film samples (Figure 1a, see the detailed methods description in Supplementary Methods). We do not generate new experimental data for this contribution. Instead, we simulate BO loops utilizing experimental BO results from article [3] and sample from the environmental stability landscape (Figure 1b) produced in [3].

The composition of a ternary perovskite $\mathrm{Cs}_x\mathrm{MA}_y\mathrm{FA}_z\mathrm{PbI}_3$, where $x + y + z = 1$, was in the previous work [3] optimized to survive environmental stress from heat, humidity, and visible light illumination. At compositions suggested by BO, perovskite thin film samples were prepared and exposed to aging tests (detailed in Supplementary Methods). Since the investigated perovskite compositions undergo color changes from dark to yellow film when degrading [14], the chosen BO target variable was the color change of the perovskite film integrated over time (during an aging test of fixed duration). The BO input variables to be optimized were the compositional proportions of A site compounds x and y

(and $z = 1 - x - y$) in the perovskite ABX_3 crystal structure. The loop of preparing samples with compositions suggested by BO (a batch-mode BO implementation with local penalization, expected improvement acquisition function, and GP surrogate model) was continued until the search converged. The GP model of environmental stability (Figure 1b) is reproduced from the converged BO search in article [3]. It should be mentioned that in [3], the expected improvement acquisition function was modified with a data fusion method to account for a secondary design requirement, thermodynamic phase stability of the compositions. In this work, we utilize the data fusion method for adding a human into the loop, but do not use the thermodynamic phase stability data.

An experienced scientist can provide an initial evaluation of perovskite film quality based on their color and uniformity. All the samples produced were initially photographed in [3]. Thus, in this work, we graded the sample quality on a scale of 0 (good quality) - 3 (low quality) post-hoc based on the saved photographs (example pictures shown in Supplementary Figure 1). We subsequently trained a GP model on this data (Figure 1c, see Supplementary Methods for details) that serves as human opinion in the present study. Interestingly, the resulting human opinion model (formed with no knowledge on the individual sample compositions) closely resembles the thermodynamic phase stability approximation in [3] (see Supplementary Figure S2). This suggests human evaluation on perovskite film samples indeed links to physically meaningful properties of the samples.

The comparison of the two GPs (Figure 1b-c) reveals a global optimum region (Region A) in the search space that has good quality samples and is environmentally stable. There is also a false optimum region (Region B): it results in seemingly almost equally good environmental stability but many of the samples are actually low quality, leading to invalid stability measurements in this region (the reasons discussed in Supplementary Section A.2). Hence, in this work, the focus is on investigating if human input can stir the BO search away from Region B and facilitate convergence towards the optimum in Region A.

2.2 Human-in-the-loop implementation

Human integration requires two steps: deciding when human opinion is required and integrating it into the BO loop (Figure 1a, details in Supplementary Methods). For both steps, we have many choices. This work is not a comprehensive evaluation of all possible choices, but instead serves as a proof-of-concept human-in-the-loop configuration that provides promising results for materials optimization loops.

Here, we request human evaluation on the next samples suggested by BO only when there are no previous human evaluations on the compositions nearby (the radius of the compositional exclusion circle serving effectively as a hyperparameter) or when there are rapid changes in the environmental stability as a function of composition (the limit value of the gradient serving as another hyperparameter). The hyperparameter values are chosen based on experimental domain knowledge (see Supplementary Methods section).

The human evaluations are, in this work, integrated into the BO loop via data fusion [3] (inspired by the unknown constraint method [15]): the point-wise human opinions are approximated with a GP surrogate model that spans the whole search space. The surrogate model is scaled sigmoidally to a probability distribution (the scaling factor serving effectively as a hyperparameter) and integrated into the BO loop in the form of an added cost to the acquisition function.

2.3 Convergence results

We evaluate the success of the human integration based on three criteria: the optimum environmental stability value found by BO, the compositional distance of this composition from Region A (evaluated as the simple regret), and the number of queries made on the human (here, number of samples from the human evaluation model). The BO loop is initiated with two randomly selected samples and there are 100 rounds in the whole BO cycle. To gather statistics, the BO cycle is repeated 25 times.

Our results show that human integration leads to faster convergence of the search to Region A (Figure 2a): the human-aided version converges in 15 rounds of sampling to a lower regret than the vanilla BO does during the whole cycle of 100 rounds. These results are gained by showing on average 5% of the samples to the human (the development shown in Figure 2b). It should be noted that whether the algorithm converges into Region A or Region B cannot be detected from the convergence plot

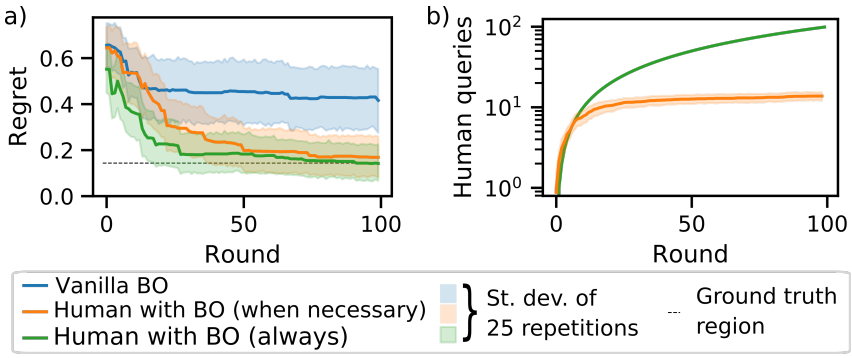

Figure 2: Results of the simulated Bayesian optimization (BO) loops with human or without (vanilla BO). The BO loop is repeated 25 times for statistics, 100 rounds of sampling is performed in each loop, and one sample is sampled in each round. a) The best environmental stability value found by the BO via a proxy variable of integrated color change during aging test duration, I_c. b) Simple regret i.e. the compositional distance of the sample with the best environmental stability found by the BO from the ground truth optimum. c) Accumulated number of queries made on the human.

only (Supplementary Figure S3). This is because both regions seem good based on the stability proxy – however, this result is invalid in Region B due to low quality samples, which is why human opinions are valuable here. The ideal rate of showing samples to a human depends on the preferred balance of increasing trust on the BO results (in which case a lot of samples should be shown to a human) or facilitating accelerated BO cycle (less samples should be shown to a human). In the low extremum, no samples are shown to a human and the performance corresponds to the vanilla BO case in Figure 2. In the high extremum, showing all the samples to a human results only in marginal added convergence benefits (Figure 2b) with a lot more effort invested into the evaluations (Figure 2c).

Adding human into the BO loop affects the acquisition locations, which indirectly introduces bias from human opinions into the resulting surrogate model of the target property, here environmental stability. Therefore it is specifically important to analyze the results of the human-in-the-loop BO in conjunction with the mean and variance predictions from the surrogate models of both the target property and human evaluation. The predictions from the surrogate model are less reliable in regions that have been sampled sparsely, thus regions avoided due to the sample quality issues need to be considered for future searches with refined sample preparation methods. The human-in-the-loop implementation in this work does not remove the samples recognized as low-quality from the target property dataset. This choice was made in order to make it easier to follow BO progress with accumulating samples. Future investigations are needed to evaluate, if this choice should be refined. Future work will also investigate different human-in-the-loop techniques and implementations.

Our demonstration is on perovskite thin film data, for which visual inspection of sample quality was a suitable way to quickly estimate sample quality. Sample quality issues are present also in other types of materials samples and similar approaches for pre-screening sample quality by an experienced human scientist can be found for many of them. Thus, the presented approach is applicable also beyond perovskite optimization tasks.

3 Conclusions

Our proof-of-concept human-in-the-loop Bayesian optimization process demonstrates that adding human into experimental materials optimization loops for evaluating sample quality can facilitate the convergence of Bayesian optimization into meaningful optima with samples of sufficient quality (in contrast to wasting sampling to regions with low-quality samples and, consequently, invalid data). Here, we have shown an example case of perovskite composition optimization for environmental stability. Our work exemplifies how the human-in-the-loop methods could be useful in AI-guided design of materials. This work aims to facilitate the movement towards increasingly high-dimensional and novel materials searches in today's partially self-driving laboratories — which are currently hindered by concerns on varying quality of the samples — and triggering further discussions on how sample quality should in future be optimally ensured in materials searches.

References

[1] Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. Journal of Global optimization, 13(4), 455-492.

[2] Greenhill, S., Rana, S., Gupta, S., Vellanki, P., & Venkatesh, S. (2020). Bayesian optimization for adaptive experimental design: a review. IEEE access, 8, 13937-13948.

[3] Sun, S., Tiihonen, A., Oviedo, F., Liu, Z., Thapa, J., Zhao, Y., ... & Buonassisi, T. (2021). A data fusion approach to optimize compositional stability of halide perovskites. Matter, 4(4), 1305-1322.

[4] Löfgren, J., Tarasov, D., Koitto, T., Rinke, P., Balakshin, M., & Todorović, M. (2022). Machine Learning Optimization of Lignin Properties in Green Biorefineries. ACS Sustainable Chemistry Engineering, 10(29), 9469-9479.

[5] Wakabayashi, Y. K., Otsuka, T., Krockenberger, Y., Sawada, H., Taniyasu, Y., & Yamamoto, H. (2019). Machine-learning-assisted thin-film growth: Bayesian optimization in molecular beam epitaxy of SrRuO3 thin films. APL Materials, 7(10), 101114.

[6] Langner, S., Häse, F., Perea, J. D., Stubhan, T., Hauch, J., Roch, L. M., ... & Brabec, C. J. (2020). Beyond ternary OPV: high-throughput experimentation and self-driving laboratories optimize multicomponent systems. Advanced Materials, 32(14), 1907801.

[7] Mikkola, P., Martin, O. A., Chandramouli, S., Hartmann, M., Pla, O. A., Thomas, O., ... & Klami, A. (2021). Prior knowledge elicitation: The past, present, and future. arXiv preprint arXiv:2112.01380.

[8] Daee, P., Peltola, T., Soare, M., & Kaski, S. (2017). Knowledge elicitation via sequential probabilistic inference for high-dimensional prediction. Machine Learning, 106(9), 1599-1620.

[9] Sundin, I., Peltola, T., Micallef, L., Afrabandpey, H., Soare, M., Mamun Majumder, M., ... & Kaski, S. (2018). Improving genomics-based predictions for precision medicine through active elicitation of expert knowledge. Bioinformatics, 34(13), i395-i403.

[10] Bharti, A., Filstroff, L., & Kaski, S. (2022). Approximate Bayesian Computation with Domain Expert in the Loop. arXiv preprint arXiv:2201.12090.

[11] Hvarfner, C., Stoll, D., Souza, A., Lindauer, M., Hutter, F., & Nardi, L. (2022). π BO: Augmenting Acquisition Functions with User Beliefs for Bayesian Optimization. arXiv preprint arXiv:2204.11051.

[12] Mikkola, P., Todorović, M., Järvi, J., Rinke, P., & Kaski, S. (2020, November). Projective preferential Bayesian optimization. In International Conference on Machine Learning (pp. 6884-6892). PMLR.

[13] Liu, Z., Rolston, N., Flick, A. C., Colburn, T. W., Ren, Z., Dauskardt, R. H., & Buonassisi, T. (2022). Machine learning with knowledge constraints for process optimization of open-air perovskite solar cell manufacturing. Joule, 6(4), 834-849.

[14] Keesey, R., Tiihonen, A., Siemenn, A. E., Colburn, T. W., Sun, S., Hartono, N. T. P., ... & Buonassisi, T. (2022). An Open-Source Environmental Chamber for Materials-Stability Testing Using an Optical Proxy.

[15] Gelbart, M. A., Snoek, J., & Adams, R. P. (2014). Bayesian optimization with unknown constraints. arXiv preprint arXiv:1403.5607.