

ARTICLE OPEN



Machine-learned metrics for predicting the likelihood of success in materials discovery

Yoolhee Kim ¹✉, Edward Kim¹, Erin Antono¹, Bryce Meredig¹ and Julia Ling¹

Materials discovery is often compared to the challenge of finding a needle in a haystack. While much work has focused on accurately predicting the properties of candidate materials with machine learning (ML), which amounts to evaluating whether a given candidate is a piece of straw or a needle, less attention has been paid to a critical question: are we searching in the right haystack? We refer to the haystack as the design space for a particular materials discovery problem (i.e., the set of possible candidate materials to synthesize), and thus frame this question as one of design space selection. In this paper, we introduce two metrics, the predicted fraction of improved candidates (PFIC), and the cumulative maximum likelihood of improvement (CMLI), which we demonstrate can identify discovery-rich and discovery-poor design spaces, respectively. A combined classification system, composed of the CMLI and PFIC metrics, is then used to identify optimal design spaces with high precision, and thus show the potential to significantly accelerate ML-driven materials discovery.

npj Computational Materials (2020)6:131; <https://doi.org/10.1038/s41524-020-00401-8>

INTRODUCTION

A significant challenge in materials discovery is the vast, often untenable, space of potential experiments that could be performed in any given materials optimization effort. Desirable, novel materials exist as needles in enormous proverbial haystacks. Brute force searches of these haystacks, which may represent material compositions, crystal structures, or synthesis parameters, are prohibitively expensive and time-consuming. Therefore, efficient methods for discovering needles are sought to reduce the number of experiments required to discover novel materials that meet a given set of performance specifications^{1–4}.

Computational techniques may alleviate this challenge by screening promising candidate materials via predicting whether each candidate from the haystack exceeds some performance threshold (e.g., thermoelectric figure of merit). This haystack may represent a large, high-dimensional design space^{1,2,5–8}, or the set of all relevant materials whose performances are unknown.

Model accuracy is a well-studied factor in determining whether a materials discovery project using computational techniques will be successful^{9–11}. Model accuracy has therefore been widely used as a success metric for machine-learned structure–property relations. A less widely considered factor is the quality of the haystack, or the design space. Depending on the quality of the design space, discovering an improved material may range from trivially easy (if all untested materials are superior) to impossible (if all untested materials are inferior). Jia et al.¹² and Kauwe et al.¹³ have shown, respectively, that random and ML-guided materials search strategies may achieve success depending on the underlying difficulty—or quality—of the design space. Thus, there is a critical need for a quantitative method that determines the quality of the design space, which strongly impacts materials discovery success. In other words, we desire a method for finding the best haystacks in which to search for needles.

To the best of the authors' knowledge, there is no current work on quantifying the quality of a given design space, where a high-quality space yields materials discovery success with fewer required experiments¹⁴. While recent studies have begun to

explore the effect of choosing different training and test sets for discovery on the practicality of discovering new and useful materials, broadly quantifying these effects in terms of successful materials discoveries remains unexplored^{12,13,15,16}.

In this work, we first demonstrate how design space quality is a critical factor in materials development success. We illustrate this by benchmarking sequential learning success versus design space quality through simulated sequential learning on existing materials datasets. Second, we introduce a procedure to initialize the training data and design space, which better reflects how known and unknown materials data are distributed^{10,13,15}. Finally, we present a design space evaluation system using two novel metrics—the PFIC and CMLI scores—for quantifying the quality of a given design space. This design space evaluation system gives insight into the likelihood of success in a given materials development project, and enables data-driven selection between different materials discovery projects.

RESULTS

Effect of design space quality on sequential learning success

Sequential learning, also known as active learning, is a data-driven workflow used to accelerate materials discovery projects by reducing the number of experiments required to find a material that meets a given set of target specifications. This workflow involves pairing a machine-learning model with an optimization algorithm to search over the design space for the most promising candidate materials, performing experiments on these candidates suggested by the model, and then retraining the model on the results of these experiments so that the model becomes better informed after each iteration¹⁷.

It would be extremely valuable to a researcher to know from the outset of a materials discovery project whether the project is likely to take 10 or 100 experiments to find an improved material. For this reason, this paper quantifies sequential learning success as the number of iterations required to find an improved candidate in the design space (i.e., the number of draws from the haystack

¹Citrine Informatics, Redwood City, CA 94063, USA. ✉email: ykim@citrine.io

Table 1. Benchmark datasets used in sequential learning simulations and analysis of predictive design space metrics.

Dataset	Origin	Dataset size	Output property	Objective
Harvard Clean Energy Project ²⁵	DFT (Density Functional Theory)	30,000	Power conversion efficiency	Maximize
Materials Project ¹⁸	DFT	30,000	HOMO energy	Maximize
			Formation energy	Minimize
			Bandgap	Tune to 1.5 eV
Melting Points ²⁶	Experiment	3041	Melting-point temperature	Maximize
Superconductors ²⁷	Experiment	585	Superconducting T_c	Maximize
UCSB Thermoelectrics ²⁸	Experiment	1093	Seebeck coefficient	Maximize
			Electrical resistivity	Minimize
Strehlow & Cook ²⁹	Experiment	1449	Bandgap	Tune to 1.5 eV

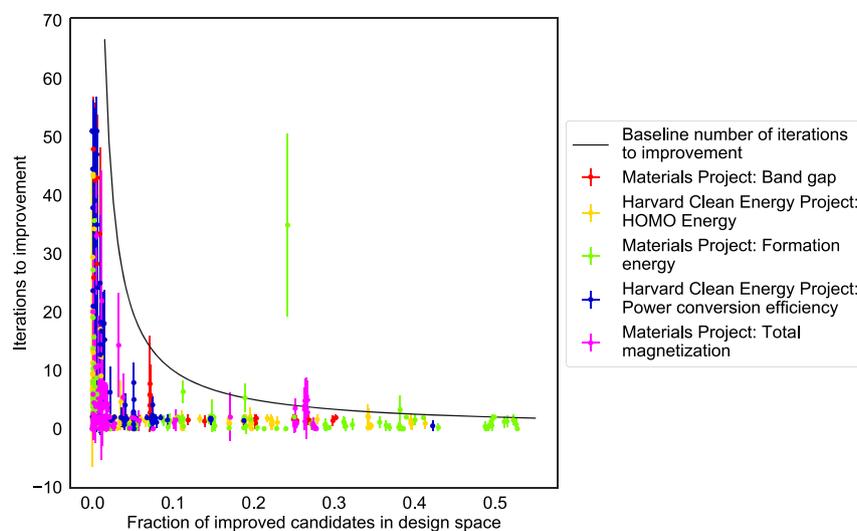


Fig. 1 Number of iterations required to find an improved candidate versus the fraction of improved candidates (FIC) in the design space. Sequential learning simulations were run over five different benchmark simulations, represented by the different colors. Each point represents a different design space and training data initialization, and the error bars on each point represent the variance in the number of iterations it takes to find an improved material over 20 sequential learning simulations with the same data initialization but stochastically trained random forest models. The black line represents the expected number of iterations to find an improved candidate via random search strategy, $1/p$, where p is the FIC, or the fixed per-trial probability of success.

until a needle is found). This paper quantifies design space quality as the fraction of improved candidates (FIC), or the fraction of candidates in the design space that perform better than the best training candidate (i.e., the fraction of the haystack that are needles). In practice, design space quality is unknown when searching for an undiscovered material. Therefore, predictive metrics correlated with the FIC are useful to assess design space quality.

In order to determine the effect of design space quality on sequential learning success, sequential learning simulations described in the Methods section were run using the Materials Project and Harvard Clean Energy Project datasets from Table 1. These datasets were subsampled and divided into a training set and a design space using the data initialization method described in the Methods section. The five properties used for these simulations, as well as their objectives, are listed in Table 1.

Figure 1 shows sequential learning success, or the number of iterations required to find an improved candidate from the design space, as a function of the FIC, or the fraction of improved candidates in the design space, over these benchmark test cases.

Each color in this scatterplot corresponds to a different sequential learning objective—maximizing, minimizing, or tuning of some target property from a benchmark dataset. Each data point in this scatterplot represents a separate data initialization of the training data and design space, and the error bars represent the standard deviation over 20 trials of identical data initializations caused by stochastic model training.

This figure shows that the number of iterations required to find an improved candidate is highly sensitive to the design space quality, or the FIC. For design spaces with low FIC, many iterations are required on average to find an improved candidate, and the number of iterations required has high variance. For design spaces with high FIC, a comparatively lower number of iterations are required to find an improved candidate. This result emphasizes the contribution of the design space quality to the difficulty of a sequential learning problem. Therefore, design space quality is expected to have a strong impact on the success of sequential learning projects, and materials discovery projects in general. Additional experiments were performed to analyze the relationship between model quality and the required number of iterations

until an improvement is found. The results of these studies indicate that model quality, unlike design space quality, is not as strongly tied to sequential learning success. The details of these tests along with their results are available in the Supplementary Discussion and Supplementary Figs. 1–6.

In addition, Fig. 1 highlights how sequential learning nearly always outperforms the baseline random-selection strategy, which is shown by the black curve. Assuming that design space candidates are chosen randomly at each sequential learning iteration, the required number of iterations to find an improved candidate can be approximated as the expected number of trials until success for a Bernoulli process. This expected number of trials is given by $1/p$, where p is the fixed per-trial probability of success, which is equal to the FIC. This baseline curve is an approximation and is only accurate for large design spaces where sampling without replacement does not significantly affect the probability of success p , or the probability for finding an improved candidate. We recognize that random selection is not representative of a typical experimental design process; however, recent work¹² has shown that random design can be just as likely to yield successful outcomes in chemical synthesis when compared to human-guided approaches.

Figure 1 also shows a clear outlier corresponding to the formation energy from the Materials Project dataset¹⁸. Upon investigating this outlier case, it was found that the number of iterations to improvement shows a bimodal distribution. In some trials, the model is able to find an improvement immediately, in the first iteration of sequential learning. In other trials, the iterations to improvement are much larger, between 21 and 46 iterations. Based on this evidence, our hypothesis is that there are several poor-performance candidates near the improved candidates in the design space. The model may get “lucky” and discover an improved material immediately, or get “unlucky” and discover the bad candidate, misleading the model to not continue exploring that region near the other improved candidates. This outlier highlights the existence, and importance, of antagonistic data splits in the validation and application of machine-learning models for materials discovery. Future work includes further investigation of this effect, and determining how to mitigate this outcome in experimental (nonsimulated) sequential learning applications.

In summary, the fraction of improved candidates in the design space (FIC) was identified as a key metric that is strongly correlated to how many experiments are required to find an improved material (i.e., the number of draws from the haystack until a needle is found), emphasizing the contribution of design space quality to the difficulty of a sequential learning project. The FIC can therefore be used to determine how likely a materials discovery project is to succeed. However, we emphasize that in general, we do not know the FIC for a given design space a priori. Thus, we would like to identify other metrics we can calculate based only on the initial training data that are predictive of the FIC.

Predictive metrics to evaluate design space quality

Given the strong relationship between design space quality and sequential learning success (the number of iterations required to find an improved candidate), design space quality can be used as a key indicator to determine the difficulty of a sequential learning problem. However during in-lab materials discovery, design space quality, or the FIC (fraction of improved candidates in the design space), is not known. Therefore, predictive design space metrics to assess the design space quality are desirable. In other words, we would like metrics that evaluate how difficult it is to find a needle in a given haystack.

Several predictive design space metrics were examined in this work. Each metric leveraged machine-learned predictions, and

Table 2. Correlations of predictive metrics with true design space quality.

Predictive metric	Pearson correlation with true FIC
PFIC	0.38
PFIC minus 1σ uncertainty	0.20
Top-10 cumulative MLI (CMLI)	0.31
Top-5 CMLI	0.27
Predicted fraction of improved, interpolated candidates	0.11
Ratio of extrapolated to all design space candidates	0.29

PFIC predicted fraction of improved candidates, FIC design space quality, CMLI cumulative maximum likelihood of improvement.

Predictive metrics are tested across the datasets in Table 1 with ten trials for each data initialization. Metrics in boldface yielded the highest correlations with true design space quality.

some of these metrics leveraged uncertainty estimates. These metrics were then compared against the FIC, the true design space quality. These metrics were calculated for ten properties over six different benchmark datasets from Table 1, which represent different materials classes, as well as both computational and experimental datasets. Each dataset had 50–100 different training data and design space initializations, where each split had varying FIC, or design space quality. Critically, sequential learning iterations were not required to calculate these predictive design space metrics and the FIC, since these quantities can be computed without any iterative data acquisition. Table 2 presents the Pearson correlations of the examined predictive design space metrics against the true design space quality, or the FIC. Additional details regarding the testing of these predictive metrics, along with their explicit definitions, can be found in the Supplementary Methods, Supplementary Table 2, and Supplementary Figs. 7–14.

As shown in Table 2, the PFIC and top-10 CMLI score achieved the highest correlations with true design space quality. Consequently, these two predictive metrics are further discussed in the following sections of this work.

PFIC score: The PFIC score is defined as the fraction of design space candidates that are predicted by the model to have improved performance over the best training data point. A machine-learning model was fit to the training data to calculate the PFIC score. After training this random forest, the model then predicted the performance of all of the candidates in the design space. The PFIC score can then be calculated by

$$\text{PFIC} = \frac{N_{p(x_i) > b}}{|X|} \quad (1)$$

In Eq. (1), X is the design space, x_i represents each candidate in the design space, $p(x_i)$ represents the predicted performance of that candidate, b is the performance of the best training data point, $|X|$ is the design space size, and N_a is the number of candidates satisfying condition a . Therefore, the numerator of this equation represents the number of design space candidates predicted to be improvements over the best training data point, and the denominator represents the design space size. Equation (1) assumes a maximization objective. A minimization objective can be achieved by subtracting the PFIC score from 1.

This work used a random forest as the machine-learning algorithm¹⁹, where a linear model was fit to the training data at each leaf node (see the Methods section for more model details). Fitting a linear model at each leaf node of the random forest was a critical choice for calculating the PFIC score. In typical random forest implementations, each leaf predicts a constant value equal

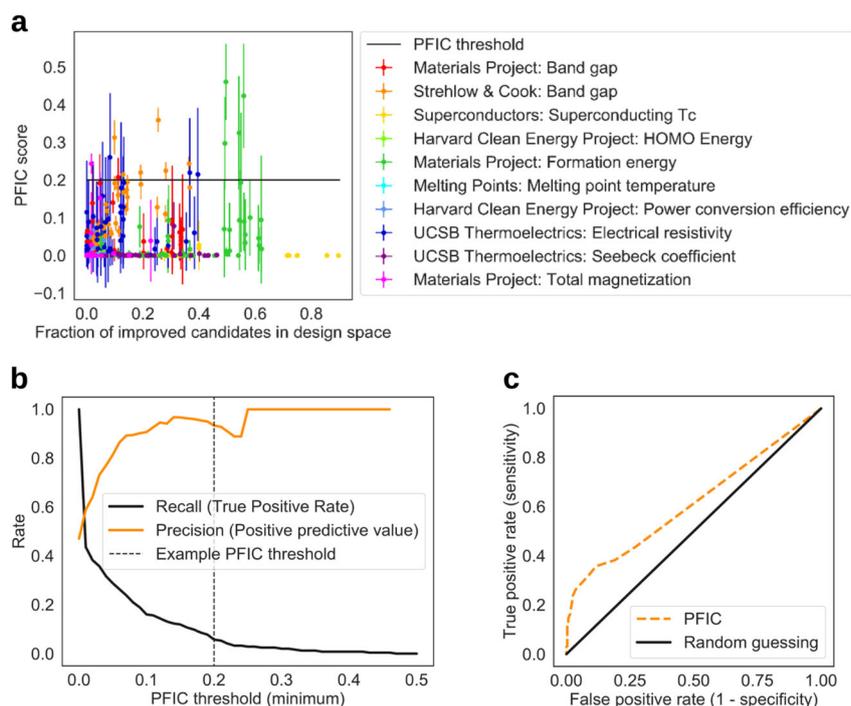


Fig. 2 PFIC score. **a** PFIC score evaluated over benchmark design spaces versus the actual fraction of improved candidates in the design space. Each point represents a different design space and training data initialization, and the error bars on each point represent the standard deviation in the PFIC score over 20 identical data initializations. The black line indicates an example threshold value of $t_{\text{PFIC}} = 0.2$, which is later used in a combined design space evaluation system. **b** Precision and recall of the PFIC score as a function of the PFIC threshold, t_{PFIC} , in identifying high-quality design spaces, where high quality is defined as containing at least 4% improved candidates in the design space. This 4% corresponds to a baseline of 25 iterations, until an improvement is found via random search over the design space. The dashed black line indicates an example threshold value of $t_{\text{PFIC}} = 0.2$. **c** ROC (receiver operating characteristic) curve for the PFIC score in identifying design spaces with a fraction of improved candidates $>4\%$. The AUC score for this curve is 0.62.

to the average of the values of the training data at that leaf. Random forests that use this averaging approach are unable to predict values outside the range of the training data, so no candidates could be predicted to be an improvement over the best training data point. Other regression models that can extrapolate beyond the range of the training data to predict improvements could be paired with the PFIC score. Such regression models include Gaussian process regressors, support vector regressors, neural networks, kernel ridge regressors, and polynomial regressors. While these alternative algorithms were not investigated in this study, they would be of interest for future work.

The PFIC score was tested on ten different benchmarks from Table 1. The results of these tests are shown in Fig. 2a. This figure shows the PFIC score as a function of the FIC, the true design space quality, across the benchmark datasets. Again, each point represents a separate data initialization, and the error bars represent the standard deviation over 20 trials of the same data initialization caused by stochastic model training. These results demonstrate that the PFIC score can be used to surface design spaces with high FIC at the outset of a project. In other words, the PFIC score could be used to highlight sequential learning projects that are likely to succeed because the haystack is full of needles. While this score does not correctly surface all high-quality design spaces (i.e., there are some false negatives), it is still able to identify many high-quality design spaces.

Figure 2b shows the precision and recall of the PFIC score in identifying high-quality design spaces as a function of the PFIC threshold, t_{PFIC} . In this figure, we defined a high-quality design space as containing at least 4% improved candidates in the design space. This 4% corresponds to a baseline of 25 iterations, until an improvement is found via random search over the design space,

which is within the typical range for sequential learning applications^{20,21}. Precision is given by $\frac{TP}{TP+FP}$, and recall is given by $\frac{TP}{TP+FN}$, where TP represents the true positives, or the design spaces correctly categorized as high quality, FP represents the false positives, or the design spaces incorrectly categorized as high quality, and FN represents the false negatives, or the high-quality design spaces incorrectly categorized as low quality. Therefore, the precision represents the fraction of the design spaces identified as high quality by the PFIC score that are actually high quality, while recall represents the fraction of all high-quality design spaces that were surfaced as high quality by the PFIC score. Fig. 2b shows that as t_{PFIC} increases, the precision improves while recall worsens. The notch in precision between $t_{\text{PFIC}} = 0.2$ and $t_{\text{PFIC}} = 0.3$ is caused by one sequential learning trial with FIC < 0.04 and PFIC > 0.2 . Between $t_{\text{PFIC}} = 0.2$ and $t_{\text{PFIC}} = 0.25$, the precision continues to decrease, as other sequential learning trials with FIC > 0.04 fall below t_{PFIC} . The precision shoots up to 1.0 once the one sequential learning trial with $t_{\text{PFIC}} < 0.04$ finally falls below t_{PFIC} at around $t_{\text{PFIC}} = 0.25$.

Figure 2c shows the ROC curve for the PFIC score in identifying high-quality design spaces. The ROC curve plots the recall against the false-positive rate, which is given by $\frac{FP}{TN+FP}$, where FP represents the false positives, or the design spaces incorrectly categorized as high quality, and TN represents the true negatives, or the low-quality design spaces correctly categorized as low quality. The AUC score for this curve is 0.62. A perfect classifier has an AUC score of 1.0, while a classifier that is not capable of distinguishing between the high- and low-quality design spaces has an AUC score of 0.5. The PFIC score thus shows some potential for the ability to surface materials discovery problems with high-quality design spaces.

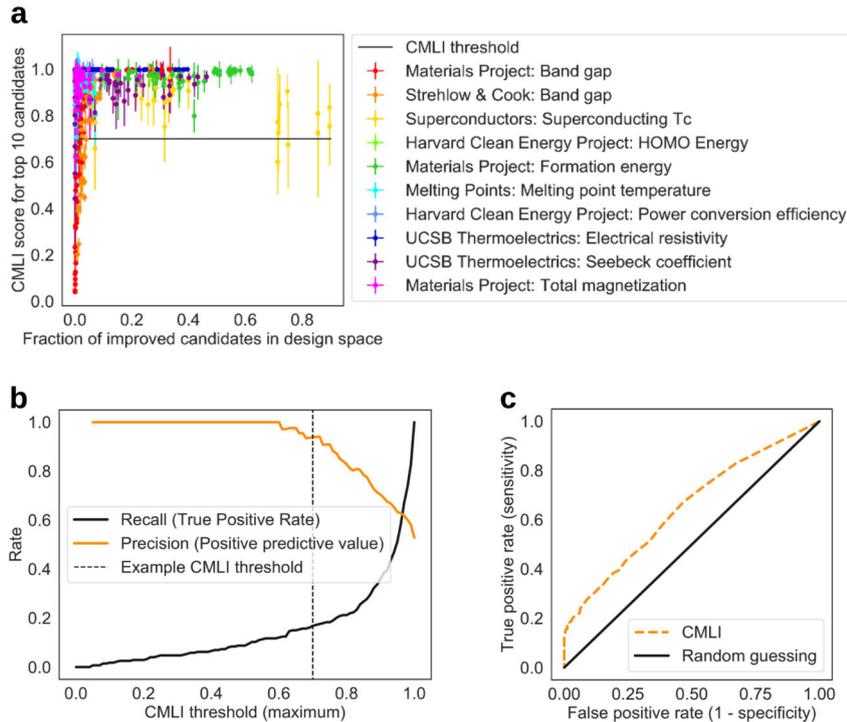


Fig. 3 Cumulative maximum likelihood of improvement (CMLI) score. **a** CMLI score for top-ten candidates evaluated over benchmark design spaces versus the actual fraction of improved candidates in the design space. Each point represents a different design space and training data initialization, and the error bars on each point represent the standard deviation in the CMLI score over 20 identical data initializations. The black line indicates an example threshold value of $t_{\text{CMLI}} = 0.7$, which is later used in a combined design space evaluation system. **b** Precision and recall of the CMLI score as a function of the CMLI threshold, t_{CMLI} , in identifying low-quality design spaces, where low quality is defined as containing <4% improved candidates in the design space. This 4% corresponds to a baseline of 25 iterations, until an improvement is found via random search over the design space. The dashed black line indicates an example threshold value of $t_{\text{CMLI}} = 0.7$. **c** ROC curve of the CMLI score in identifying design spaces with a fraction of improved candidates less than 4%. The AUC score for this curve is 0.65.

CMLI score: The CMLI score is defined as the predicted probability that at least one candidate, out of the top n candidates in the design space with the highest likelihoods of improvement, performs better than the best training data point, where n is tunable by the experimenter. Calculating this score required using a machine-learning model to predict the performance and estimate the uncertainty for each design space candidate.

This study used a random forest as the underlying machine-learning algorithm¹⁹(see the Methods section for more model details). The uncertainty estimates were calculated using a combination of jackknife-based methods and an explicit bias model^{17,22}. However, other algorithms such as Gaussian process regressors could also be used to compute the CMLI score.

First, the top n candidates with the highest likelihoods of improvement must be identified. The likelihood of improvement L of a given design space candidate x_i is given as

$$L(x_i) = \int_b^{\infty} N(p(x_i), \sigma(x_i)) dx \quad (2)$$

In Eq. (2), $N(\mu, \sigma)$ is a normal distribution, with mean μ and standard deviation σ . The mean μ is given by the predicted performance of candidate x_i , by the machine-learning model, and the standard deviation σ is given by the estimated uncertainty. Therefore, this integral represents the probability that design space candidate x_i is an improvement over the best training data point. Equation (2) assumes a maximization objective. Calculating the likelihood of improvement L for a minimization objective would require changing the integration limits to $-\infty$ and b . While normally distributed uncertainties were used in this study, L could

be readily computed for other uncertainty distributions as long as they are integrable.

After the top n candidates are identified, the CMLI score can be computed by the following equation:

$$\text{CMLI} = 1 - \prod_{i=1}^n (1 - L(x_i)) \quad (3)$$

Equation (3) takes the product of the likelihoods that candidates x_i are *not* an improvement, $1 - L(x_i)$. This product is over the top n candidates from the design space with the highest likelihood L of performing better than the best training data point. This equation assumes that these likelihoods are independent for the top n candidates. Therefore, this score assesses the likelihood that at least one of these top n candidates is an improvement. In many cases, the model predictions are not independent, so the CMLI score has a tendency to overestimate the true design space quality. As a result, using $n = |X|$ results in CMLI scores near unity for large design spaces. This study uses the CMLI score with n at a value of 10. However, we have tested other values of n , which did not substantially affect the benchmarked test results.

The CMLI score was tested on ten benchmark datasets from Table 1. The results from these tests are shown in Fig. 3a. This figure shows the CMLI score versus the actual fraction of improved candidates across these benchmark datasets. This figure shows that the benchmark cases with low CMLI scores generally correspond to lower FICs, and that the CMLI score can be used to flag low-quality design spaces at the outset of a project. Therefore, this score can be used to filter out sequential learning projects that may be difficult due to scarcity of needles in the

haystack. While this score is not able to flag all low-quality design spaces, it is able to identify some design spaces with low FIC.

Figure 3b shows the precision and recall of the CMLI score in identifying low-quality design spaces as a function of the CMLI threshold, t_{CMLI} . In this figure, we defined a low-quality design space as containing <4% improved candidates in the design space. For the CMLI score identifying low-quality design spaces, the precision represents the fraction of design spaces identified as low quality by the CMLI score that are actually low quality, while the recall represents the fraction of all low-quality design spaces that were surfaced as low quality by the CMLI score. Figure 3c shows the ROC curve for the CMLI score in identifying low-quality design spaces. The AUC score for this curve is 0.65, where a perfect classifier has an AUC score of 1.0, while a classifier that is not capable of distinguishing between the high- and low-quality design spaces has an AUC score of 0.5. The CMLI score can thus be used to flag materials discovery problems with low-quality design spaces.

Design space evaluation system

We have investigated two metrics, the PFIC and CMLI scores, that can identify high-quality and low-quality design spaces, respectively. Therefore, these two can be combined into a system for assessing design space quality. In the below description of our design space evaluation system, we refer to ML Models A and B. Model A must be capable of estimating uncertainty, and Model B must be able to predict output values that extrapolate outside the range of the training data, in order to predict improved design space candidate performance. In this work, we used the same model for Models A and B, but two different models could be used in principle.

- Build a machine-learning Model A with uncertainty estimates on the training data
- Use Model A to make predictions and uncertainty estimates for each design space candidate
- Use the predictions and uncertainty estimates from Model A to calculate the CMLI score for the top n candidates
- Build a machine-learning Model B on the training data
- Use this Model B to make predictions for each design space candidate
- Use the predictions from Model B to evaluate the PFIC score
- Set thresholds t_{CMLI} and t_{PFIC} based on the desired precision and recall
- Use the predictive metrics to assess design space quality

Figure 4 shows the set of design spaces as classified by this design space evaluation system with example thresholds of $t_{\text{PFIC}} = 0.2$ and $t_{\text{CMLI}} = 0.7$. The x axis is the PFIC for each design space, and the y axis is the CMLI for each design space. The thresholds indicated by the black line are tunable by the user, and represent the decision boundaries for the system.

In this example, we have defined high-quality design spaces as having >4% of improved candidates in the design space, and low-quality design spaces as having <4% of improved candidates. Again, this 4% corresponds to a baseline of 25 iterations, until an improvement is found via random search over the design space, which is within the typical range for sequential learning applications^{20,21}. Ideally, all design spaces with $\text{FIC} > 0.04$ should be classified as high quality, those with $\text{FIC} < 0.04$ should be classified as low quality, and there should not be any unknown quality design spaces. The collection of design spaces in dark purple and blue shows that most design spaces have less than 10% of improved candidates. This mirrors real materials development projects; very few projects have design spaces with many improvements, with the exception of nascent projects where not many materials have been explored. While the design space evaluation system is not perfect, Fig. 4 shows that the design

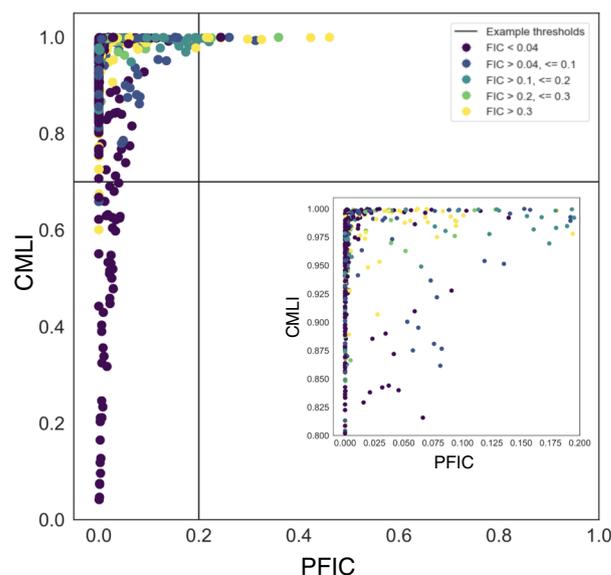


Fig. 4 Distribution of design space qualities as classified by design space evaluation system. The dashed lines indicate example thresholds of $t_{\text{PFIC}} = 0.2$ and $t_{\text{CMLI}} = 0.7$. The upper-left quadrant represents design spaces classified as unknown quality. The lower-left and upper-right quadrant show, respectively, design spaces categorized as low- and high-quality design spaces. The inset figure shows a zoomed view of the upper-left quadrant.

spaces classified as low quality are generally below 0.04 FIC, while the design spaces classified as high quality are generally above 0.04 FIC. Therefore, this system is useful in distinguishing between low- and high-quality design spaces. In other words, it is a high-precision, low-recall classifier.

With these example thresholds, we are able to identify high-quality design spaces with a precision of 0.94 and recall of 0.06. This means that 94% of the design spaces identified as high quality by the design space evaluation system are actually high quality. In addition, this means that we are able to correctly surface 6% of all the high-quality design spaces. Meanwhile, we are able to flag low-quality design spaces with a precision of 0.96 and a recall of 0.23. This means that 96% of the design spaces flagged as low quality are actually low quality, and that we are able to correctly identify 23% of all low-quality design spaces. In this example, the design space evaluation system is lossy—many design spaces are categorized as unknown quality. However, we can be fairly confident in design spaces that are categorized as high- or low quality. In addition, the thresholds for the predictive metrics can be tuned to achieve the desired precision and recall scores for both high- and low-quality design spaces.

In summary, this design space evaluation system uses both the PFIC and CMLI scores to identify high- and low-quality design spaces. We have shown that design space quality has a strong impact on simulated sequential learning projects. Furthermore, we demonstrated that this design space evaluation system is a useful tool for assessing the probability of success for any given materials development project.

DISCUSSION

While many previous studies have explored the promise of machine-learning methods for accelerating materials development, the use of machine learning to assess the difficulty of a materials development project a priori represents a new contribution. It would be extremely valuable to know how hard a materials development project would be at the outset. This information could be used to determine which projects to invest

in and how to allocate resources across a research and development portfolio. For groups investing in multiple projects simultaneously, information on project difficulty could be used to balance the project portfolio between high- and low-risk projects.

In this work, we demonstrated the importance of design space quality on materials discovery success, introduced a new data initialization method to reflect in-lab materials discovery, and finally, defined novel predictive metrics to determine the quality of a design space at the outset of a materials discovery project.

After evaluating these design space metrics across a variety of computational and experimental materials datasets, we identify two scores—PFIC and CMLI—which correlate with the true design space quality. These two metrics are then combined into a high-precision, model-agnostic design space evaluation system. Our work is thus a first step toward determining, a priori, the difficulty of a materials development project. While all the trials run in this study used a simulated sequential learning workflow, these predictive metrics could also be used in more traditional development settings where experimental test order was determined by a scientist instead of a machine-learning algorithm. In all cases, the fraction of improved candidates in the design space will be strongly related to the difficulty of a materials development project.

Future work may include testing additional predictive metrics to reveal further insights into design space quality. Future work may also include correlating predictive design space metrics directly to sequential learning success (iterations until an improvement is found) rather than the FIC, in order gain further insight into how these predictive metrics can be used to prioritize sequential learning projects. These metrics may also be calculated at every iteration of sequential learning to analyze how these metrics change when new data points are added to the training data. In addition, further experiments may be conducted to determine how model accuracy, algorithmic choice, and uncertainty quantification approach affect the accuracy of our predictive design space metrics. Finally, the metrics discussed in this work concern the optimization of single materials properties, and future work may address the challenge of multi-objective optimization.

METHODS

Overall schematic of the methodology

The methods presented in this work are highlighted in Fig. 5. Using a novel data-partitioning scheme to best represent the nature of in-lab materials discovery, a machine-learned model is trained on a set of training data. Following this, predictions for material properties (e.g., bandgap) are computed along with uncertainty estimates across materials in the held-out design space. Finally, these predictions and uncertainties are used to produce design space metrics that are indicative of the difficulty of materials discovery for a particular design space.

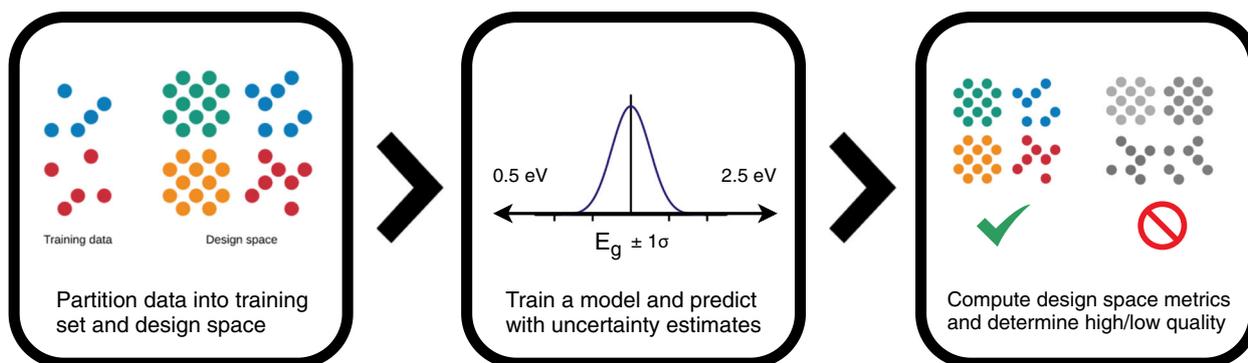


Fig. 5 Overall schematic of the methods presented in this work. Data are first initialized into a training set and a design space. This is followed by training a model using the training data, and predicting material performances across the design space candidates. Finally, predictive design space metrics are computed and used to evaluate high- and low-quality design spaces.

Toward the goals of reproducibility and open science, we have used the open-source *lolo* library in all our modeling approaches, along with a standard *k*-means clustering algorithm for data initialization applied to a collection of open-access datasets. The details of these open resources are listed throughout this section, and additional hyperparameters are detailed in the Supplementary Information.

Sequential learning simulations

To enable the rapid computation of quantitative materials discovery success and design space quality metrics across a broad variety of materials categories, as well as different splits of the training set and design space, in silico simulations of sequential learning were used to avoid the overhead of in-lab experimental trials. In silico simulations of sequential learning reduce the time required to run experiments to effectively zero.

Datasets. Sequential learning was simulated using two publicly available datasets—the Materials Project and Harvard Clean Energy Project datasets, shown in Table 1. Across a number of trials, we split each dataset into a training set and design space to assess the quality of the design space on the overall likelihood of sequential learning success. The candidates in the design spaces held out from model training were treated as an approximation to infinite design spaces commonly encountered during in-lab materials development. The methodology of splitting these benchmark datasets into a training set and design space is discussed in the following section.

The training set was used to train machine-learning models; then the trained models were used to produce predictions with uncertainties for each of the design space candidates' material properties of interest. At each sequential learning iteration, the candidate with the highest performance was then selected to “measure,” at which point the true performance of that candidate was revealed to the model, and that candidate was added to the training set. Prior to the next sequential learning iteration, the model was retrained on the training data, including the added design space candidate. Various metrics were recorded at each iteration of sequential learning for analysis, and these metrics are discussed in the following section.

Machine-learning models were built for five properties in total. The objective was to maximize, minimize, or tune the value of the target property, depending on the case. Table 1 lists the datasets and properties used for simulating sequential learning. A variety of benchmark datasets were chosen to represent different materials classes, as well as both computational and experimental datasets. The machine-learning model inputs were representative of the degrees of freedom that could be adjusted in a laboratory setting, such as the composition and the processing parameters of a material. The output of the machine-learning model was the material property to be optimized, such as the bandgap.

In Table 1, the Harvard Clean Energy Project and Melting Points datasets used SMILES (Simplified Molecular-Input Line-Entry System) strings as input. This input was featurized using the open-source software package Chemistry Development Kit (CDK)²³. This featurizer utilizes CDK's QSAR (quantitative structure–activity relationship) descriptor calculations.

The Materials Project, Superconductors, UCSB Thermoelectrics, and Strehlow & Cook datasets from Table 1 used chemical composition as

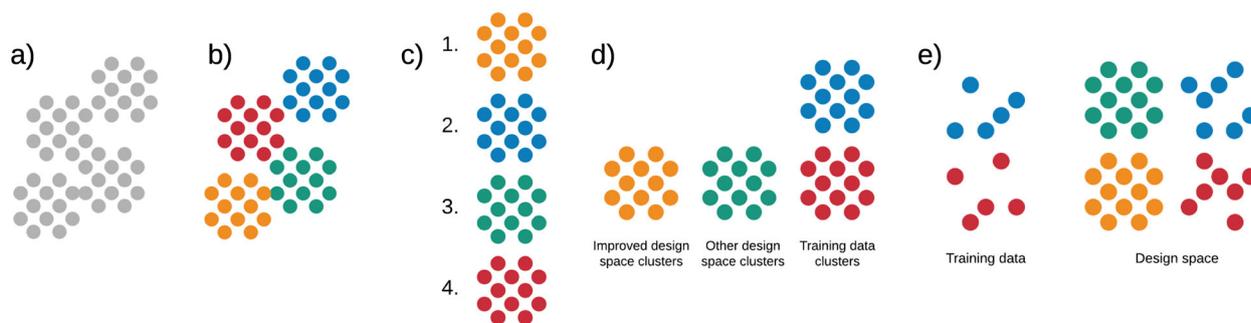


Fig. 6 Diagram of the data initialization approach. Data initialization begins with (a) a given dataset. The data initializer (b) clusters the full dataset, then (c) ranks the clusters based on the best-performing candidate. After ranking the clusters, (d) the clusters are assigned as either training clusters or design clusters. Finally, the data initializer (e) subsamples the training clusters so that some elements from the training clusters are included in the design space.

input. This input was featurized by computing elementwise means of various elemental properties taken from the software package Magpie maintained by Wolverton Research Group²⁴. This input was also featurized by calculating other compositional features, described in more detail in the Supplementary Methods.

Model details. In this work, throughout all experiments, the open-source *lolo* random forest library was used for the machine-learning algorithm¹⁹, where at each leaf node in this random forest, a linear model was fit to the training data. Minimum data per leaf were set to 20 to avoid overfitting at leaf nodes. All other hyperparameters were left at the library's default values—the number of trees was set to the training set size, the maximum tree depth was set to 30, and L2 regularization of the linear model coefficients was implemented.

Random forests make predictions for a new point based on the values of training data points that are nearby in the input space. Given the clustered nature of many materials datasets, this approach to making predictions makes random forests particularly well-suited to materials applications¹⁷. However, we stress that the methods introduced in this work are model-agnostic and are thus compatible with a wide variety of modern algorithms, including Gaussian processes and neural networks. We additionally note that all these models tend to assume that similar materials yield similar properties, and such assumptions may break down in certain situations (e.g., near a phase transition).

Model training was performed on the datasets according to the following details:

- For each training set and design space, 20 identical sequential learning trials were performed, to capture the variance due to stochastic model training.
- Each trial of the sequential learning process was performed for a maximum of 50 iterations, or until a superior candidate was discovered (where each iteration represents the addition of one new data point).
- For each dataset, 50–100 different training data and design space initializations were tested.

The maximum likelihood of improvement (MLI) and the maximum expected improvement (MEI) were used as acquisition functions to determine the performance of the design space candidates¹⁷. Each training data and design space split exhibited a varying level of sequential learning difficulty by virtue of the training-to-design-space ratio and fraction of improved candidates in the design space. Therefore, sequential learning trials were simulated in order to analyze the relationship between design space quality and success in a materials discovery project without the need for in-lab experiments.

Data initialization

While randomly chosen splits of training and test data are common throughout machine learning^{10,11}, recent works by Meredig et al. and Kauwe et al.^{13,15} have shown that the choice of training and test data significantly influences estimates of generalization error in machine-learned models, along with the success of materials discovery.

To reflect the clustered nature of materials data^{13,15}, which are generally derived from human-biased historical results¹², a *k*-means cluster-based data initialization method was developed and used in this study. This data initialization method better matched experimental materials discovery

projects, since experiments are typically run over similar types of materials, resulting in clustered training data¹². Figure 6 shows a diagram of this data initialization method, and further details regarding hyperparameter choices (e.g., number of clusters) can be found in Supplementary Table 1. We also note that a detailed discussion on the effects of data clustering and machine-learning model evaluation can be found in recent work by Meredig et al.¹⁵.

We also tested random selection of training sets and design spaces, which often yielded trivially solvable sequential learning simulations when initial training data included at least one material nearby to improved, held-out materials. Machine-learned algorithms were able to successfully discover an improved material within five iterations of sequential learning in such design spaces in almost all trials. Existing experimental and computational sequential learning studies have typically explored dozens to hundreds of iterations^{20,21} before discovering novel or improved materials, suggesting that these splits of training and test data were unrealistic, and confirming that a more complex data initialization method was required.

The data initialization method was performed via the following algorithm:

- Data were first clustered into n clusters. The number of clusters, n , in this work was tested at 16, 32, 64, and 128.
- The n clusters were ranked based on the best-performing candidate within each cluster, with respect to the material property of interest.
- These clusters were divided into training and design space clusters. Different trials varied the fraction of improved clusters that were assigned as design clusters, thereby varying the quality of the design space.
- Training clusters were split into two parts. Some of the data were included in the training set, and others were included in the design space. This approach was adopted because typical design spaces include both interpolative (similar to training data) and extrapolative (different from training data) candidates. No data points were included in both the training set and design space.

The data initialization algorithm exposed several parameters that were used to modulate the difficulty of the sequential learning problem. These parameters are discussed in more detail in Supplementary Table 1. A higher fraction of improved clusters in the design space (i.e., more needles in the haystack) corresponds to an easier materials discovery problem. In addition, the difficulty of the problem could be influenced by which points from a training cluster were included in the design space. If the best points in a given training cluster were assigned to the design space, then the materials discovery problem was found to be substantially easier because improved candidates existed near the training data. On the other hand, if no improved candidates from training clusters were included in the design space, then it was a more difficult sequential learning problem because extrapolation was required to identify improved candidates.

This clustering approach to divide the training set and design space was developed to use in our sequential learning simulations, as well as to calculate predictive design space metrics. While this data initialization approach may seem complex, it was developed with great care to realistically reflect how known and unknown data are distributed in materials datasets^{10,12,13,15}.

DATA AVAILABILITY

The datasets required to reproduce the results in this study are available on Citrination at <https://citrination.com/>. The data generated and analyzed during the study are available from the corresponding author upon reasonable request.

CODE AVAILABILITY

The codes used to generate the results in this study are Intellectual Property of Citrine Informatics, and thus are not publicly available. However, the results may be reproduced using the descriptions in this paper.

Received: 23 January 2020; Accepted: 3 August 2020;

Published online: 26 August 2020

REFERENCES

- Meredig, B. et al. Combinatorial screening for new materials in unconstrained composition space with machine learning. *Phys. Rev. B* **89**, 094104 (2014).
- Isayev, O. et al. Materials cartography: representing and mining materials space using structural and electronic fingerprints. *Chem. Mater.* **27**, 735–743 (2015).
- Rajan, K. Materials informatics: the materials "gene" and big data. *Annu. Rev. Mater. Res.* **45**, 153–169 (2015).
- Ramprasad, R., Batra, R., Piliya, G., Mannodi-Kanakkithodi, A. & Kim, C. Machine learning in materials informatics: recent applications and prospects. *npj Comput. Mater.* **3**, 54 (2017).
- Gómez-Bombarelli, R. et al. Design of efficient molecular organic light-emitting diodes by a high-throughput virtual screening and experimental approach. *Nat. Mater.* **15**, 1120 (2016).
- Greeley, J. et al. Computational high-throughput screening of electrocatalytic materials for hydrogen evolution. *Nature Mater.* **5**, 909–913 (2006).
- Kim, E., Huang, K., Jegelka, S. & Olivetti, E. Virtual screening of inorganic materials synthesis parameters with deep learning. *npj Comput. Mater.* **3**, 53 (2017).
- Borboudakis, G. et al. Chemically intuited, large-scale screening of MOFs by machine learning techniques. *npj Comput. Mater.* **3**, 40 (2017).
- Snoek, J., Larochelle, H. & Adams, R. P. Practical bayesian optimization of machine learning algorithms. in *Advances in Neural Information Processing Systems*. 2951–2959 (2012).
- Ward, L. et al. Matminer: An open source toolkit for materials data mining. *Comput. Mater. Sci.* **152**, 60–69 (2018).
- Ward, L., Agrawal, A., Choudhary, A. & Wolverton, C. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Comput. Mater.* **2**, 16028 (2016).
- Jia, X. et al. Anthropogenic biases in chemical reaction data hinder exploratory inorganic synthesis. *Nature* **573**, 251–255 (2019).
- Kauwe, S., Graser, J., Murdock, R. & Sparks, T. Can machine learning find extraordinary materials? *Comput. Mater. Sci.* **174**, 109498 (2019).
- Sanchez-Lengeling, B. & Aspuru-Guzik, A. Inverse molecular design using machine learning: generative models for matter engineering. *Science* **361**, 360–365 (2018).
- Meredig, B. et al. Can machine learning identify the next high-temperature superconductor? examining extrapolation performance for materials discovery. *Mol. Syst. Des. Eng.* **3**, 819–825 (2018).
- Raccuglia, P. et al. Machine-learning-assisted materials discovery using failed experiments. *Nature* **533**, 73 (2016).
- Ling, J., Hutchinson, M., Antono, E., Paradiso, S. & Meredig, B. High-dimensional materials and process optimization using data-driven experimental design with well-calibrated uncertainty estimates. *Integr. Mater. Manuf. Innov.* **6**, 207–217 (2017).
- Jain, A. et al. Commentary: the materials project: a materials genome approach to accelerating materials innovation. *Apl. Mater.* **1**, 011002 (2013).
- Hutchinson, M. Lolo library, v2.2.5, <https://github.com/CitrineInformatics/lolo> (2018).
- Granda, J. M., Donina, L., Dragone, V., Long, D.-L. & Cronin, L. Controlling an organic synthesis robot with machine learning to search for new reactivity. *Nature* **559**, 377 (2018).
- Brandt, R. E. et al. Rapid photovoltaic device characterization through bayesian parameter estimation. *Joule* **1**, 843–856 (2017).

- Efron, B. *Model Selection Estimation and Bootstrap Smoothing* (Division of Biostatistics, Stanford University, 2012).
- Christoph Steinbeck, D. G. & Egon, W. Chemistry development kit library, v2.3, <https://cdk.github.io/> (2018).
- Ward, L. et al. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Comput. Mater.* **2**, 16028 (2016).
- Hachmann, J. et al. The harvard clean energy project: large-scale computational screening and design of organic photovoltaics on the world community grid. *J. Phys. Chem. Lett.* **2**, 2241–2251 (2011).
- Bradley, J.-C., Williams, A. & Lang, A. Jean-Claude Bradley open melting point dataset. *FigShare* **10**, m9 (2014).
- Michel, K. Superconductor critical temperatures on Citrination. (2016). <https://citrination.com/datasets/2210/>.
- Gaultois, M. W. et al. Data-driven review of thermoelectric materials: performance and resource considerations. *Chem. Mater.* **25**, 2911–2920 (2013).
- Strehlow, W. & Cook, E. Compilation of energy band gaps in elemental and binary compound semiconductors and insulators. *J. Phys. Chem. Ref. Data* **2**, 163–200 (1973).

ACKNOWLEDGEMENTS

The authors would like to thank Astha Garg for helpful discussions in brainstorming sequential learning metrics. The authors would also like to thank Suzanne Petryk for assistance in developing the framework for sequential learning simulations.

AUTHOR CONTRIBUTIONS

Y.K. and E.A. ran the sequential learning simulations and analyzed the results. E.A. created the framework for running these sequential learning simulations. Y.K., E.K., E.A., B.M., and J.L. designed the study and wrote the paper.

COMPETING INTERESTS

The authors declare that a US patent application has been filed relating to the use of a design space evaluation system for materials development. This patent application has been filed by Citrine Informatics, with all co-authors listed as co-inventors. At the time of paper publication, a patent has been awarded under the number US 10,657,300 B1.

ADDITIONAL INFORMATION

Supplementary information is available for this paper at <https://doi.org/10.1038/s41524-020-00401-8>.

Correspondence and requests for materials should be addressed to Y.K.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020