Coherent Collections of Rules Describing Exceptional Materials Identified with a Multi-Objective Optimization of Subgroups

Lucas Foppa*[1] and Matthias Scheffler[1]

[1] *The NOMAD Laboratory at the Fritz Haber Institute of the Max-Planck-Gesellschaft and IRIS-Adlershof of the Humboldt-Universität zu Berlin, Faradayweg 4-6, D-14195 Berlin, Germany*

Using a modest amount of data from a large population, subgroup discovery (SGD) identifies outstanding subsets of data with respect to a certain property of interest of that population. The SGs are described by "rules". These are constraints on key descriptive parameters that characterize the material or the environment. These parameters and constraints are obtained by maximizing a quality function that establishes a tradeoff between SG size and utility, i.e., between generality and exceptionality. The utility function measures how outstanding a SG is. However, this approach does not give a unique solution, but typically many SGs have similar quality-function values. Here, we identify coherent collections of SGs of a "Pareto region" presenting various size-utility tradeoffs and define a SG similarity measure based on the Jaccard index, which allows us to hierarchically cluster these optimal SGs. These concepts are demonstrated by learning rules that describe perovskites with high bulk modulus. We show that SGs focusing on exceptional materials exhibit a high quality-function value but do not necessarily maximize it. We compare the mean shift with the cumulative Jensen-Shannon divergence (D_{sJS}) as utility functions and show that the SG rules obtained with D_{cJS} are more focused than those obtained with the mean shift.

I. INTRODUCTION

Machine learning typically focuses on global approaches that attempt to describe all possible scenarios with a single model. This often implies that outstanding situations may be averaged away by the regularization. In materials science, this is a severe shortcoming, because the number of possible materials is practically infinite and the compounds that are good for a certain application are statistically exceptional. In contrast, subgroup discovery (SGD) [1–4] is able to concentrate on outstanding situations, accepting that the mechanisms governing the materials' performance might be different for different SGs of materials across the immense materials space.

The SGD approach is based on a data set for a subpopulation $\widetilde{P}$ of materials, which is part of a larger space of possible materials, the full population, P. Each material of P is characterized by potentially relevant, known features, hereafter *candidate descriptive parameters*. The target of interest, e.g., a certain materials property or function, is only known for the materials in $\widetilde{P}$. SGD starts by generating a pool of propositions about the candidate descriptive parameters. Each proposition is only verified for a portion of $\widetilde{P}$. For the case of continuous (metric) parameters, the propositions are typically inequalities constraining the values of the parameters to some minimum or maximum thresholds to be determined during the analysis. Then, SGD uses a search algorithm, for instance Monte-Carlo-based[5, 6] or branch-and-bound,[7] to identify selectors σ. These are statements formed by a number of propositions and the "AND" ($\wedge$) connector, that result in SGs that maximize a quality (objective) function $Q(SG, \widetilde{P})$ of the form

$$Q(SG, \widetilde{P}) = \left(\frac{s(SG)}{s(\widetilde{P})} \right)^{\alpha} (u(SG, \widetilde{P}))^{\beta}, \qquad (1)$$

where $s(SG)$ and $s(\widetilde{P})$ are the sizes of the SG and of $\widetilde{P}$, respectively, i.e., the number of data points in the SG and in $\widetilde{P}$. The ratio between the SG size and the dataset size, $s(SG)/s(\widetilde{P})$, is referred to as the *coverage*. $u(SG, \widetilde{P})$ is a utility measure (or utility function) describing how outstanding the distribution of the target in the SG is compared to the distribution of the target in $\widetilde{P}$. The utility function is chosen according to the question to be addressed, and there are many possibilities such as the shifts of mean or median in the SG with respect to $\widetilde{P}$, in the case of metric targets.[4] Finally, α and β are tunable parameters controlling the tradeoff between the SG size and its utility. Usually, $\alpha = \beta = 1$ or $\beta = 1 - \alpha$, with $\alpha \in [0.1, 0.9]$. The selectors identified in the SGD analysis depend on the key descriptive parameters that are typically few, out of all initially offered candidate parameters. In analogy to genes in biology, these key descriptive parameters might be called *materials genes*,[8] as they correlate with the materials property of interest. The propositions entering these selectors, in turn, are interpretable *rules* describing the outstanding behavior of the SG, for instance, the exceptional performance of materials. These rules can then be exploited to efficiently identify outstanding materials in P for which the target is not known.

The potential of SGD to uncover locally outstanding patterns in materials science has been demonstrated by several works.[3, 4, 9–11] In particular, the ability of SGD to handle unbalanced datasets dominated by low-performance situations, which are ubiquitous in materials science, has been highlighted.[11] However, previous

works generally focused on the single SG that maximizes the quality function for a fixed value of α. This description does not reflect all possible size-utility or, equivalently, generality-exceptionality, tradeoffs that might be relevant for a given application.[12] Besides, the definition of the utility function often assumes that the distribution of the target in $\widetilde{P}$ is appropriately represented by a single value and that this value properly reflects the larger, unknown population P. For instance, the mean value of the target, evaluated using all data points in $\widetilde{P}$, can be used in utility functions that favor the selection of SGs with higher or lower values of targets. However, in the typical scenario of materials science the assumption may be questioned, as high-quality data sets are small and the known distributions are often nonnormal or highly unbalanced.

In this manuscript, a *Pareto region* concept is introduced to identify coherent collections of SG rules that correspond to a multitude of size-utility tradeoffs. Furthermore, we discuss the similarity between SGs based on the Jaccard index. This similarity measure enables the hierarchical clustering of the identified optimal SG rules. The mean shift between the distributions of target values in the SG and in $\widetilde{P}$ is also compared with the cumulative Jensen-Shannon divergence (D_{cJS})[13] as an alternative utility function.

The concepts are demonstrated by the identification of rules that describe cubic ABO_3 perovskite materials with high bulk modulus. We show that useful rules focusing on exceptional materials do not necessarily correspond to the single SG which maximizes Q, but they can be systematically identified with the Pareto-region approach. Moreover, SGs obtained by considering the nonparametric D_{cJS} as utility function provide a more focused description compared to those obtained with the mean shift.

II. RESULTS

Pareto Region Identification and Similarity Analysis

In order to identify the pursued coherent collections of SGs with multiple size-utility tradeoffs, we first run the SGD algorithm using the quality function of Eq. 1 with $\alpha = \beta = 1$ and two different utility functions: the positive mean shift and the cumulative Jensen-Shannon divergence D_{cJS} (see below). Thus, coverage and utility are given the same importance. Then, we collect a number of SGs identified by this algorithm that are top-ranked according to Q. Among these SGs, we identify a Pareto region with respect to the two objectives coverage and utility function. In multi-objective optimization, a Pareto front is defined as the set of solutions for which no single objective can be improved without deteriorating at least one other objective. Thus, the solutions in the Pareto front reflect an optimal tradeoff between competing objectives. To ensure that no interesting SG is left

out, we included in our analysis, not only SGs that are part of the Pareto front, but also SGs within a given fixed, small distance to the Pareto front in the coverage-utility space, i.e., SG which are *near the Pareto front*. This is called the *Pareto region*.

In order to assess the variability of the SG rules within the Pareto region, we define a similarity measure between two SGs using the Jaccard similarity index J:

$$J(SG_i, SG_j) = \frac{|SG_i \cap SG_j|}{|SG_i \cup SG_j|}. \tag{2}$$

$J(SG_i, SG_j)$ corresponds to the number of data points shared by two SGs, SG_i and SG_j, divided by the number of data points corresponding to the union of the two SGs. J ranges from 0 to 1, and the higher the value, the more similar the SGs. We use the $J(SG_i, SG_j)$ values to create a *similarity matrix* containing the Jaccard indices for all combinations of SGs in the Pareto region. Obviously, $J(SG_i, SG_j) = J(SG_j, SG_i)$, i.e., this matrix is symmetric.

Finally, we group the SGs of the Pareto region according to their similarity by applying agglomerative hierarchical clustering[14] to the similarity matrix. In this unsupervised clustering approach, each data point (SG) is initially considered a cluster. Then, the two most alike clusters with respect to a linkage criteria are combined (agglomerated), forming a single bigger cluster. This procedure is repeated until all data points are contained in one large cluster. The clustering analysis results in a sorted similarity matrix, where similar SGs are close to each other. Moreover, a tree-like diagram is generated. The bottom of this diagram contains the individual SGs (tree leafs). These leafs are linked by branches that are combined as the height increases. This tree-like structure shows how the individual SGs and clusters were agglomerated at each iteration of the algorithm. It is called dendogram, and it can be truncated at any height to create an arbitrary number of clusters between 1 and the total number of SGs in the Pareto region. These clusters can then be used to analyze the coherent collection and to choose the SG rules that should be considered for further investigation of materials genes or for exploitation, i.e., for the discovery of new materials. We note that the number of clusters is not defined beforehand, but one rather chooses the desired number of cluster to consider when truncating the dendogram. Further details on the SGD approach, Pareto-region identification, and hierarchical clustering are available in the Methods section.

Application to Perovskite Design

The identification and clustering of coherent collections of SGs and the usefulness of our approach will be demonstrated for the learning of rules describing ABO_3 perovskites that exhibit a high bulk modulus, B_0. The offered candidate descriptive parameters are atomic prop-
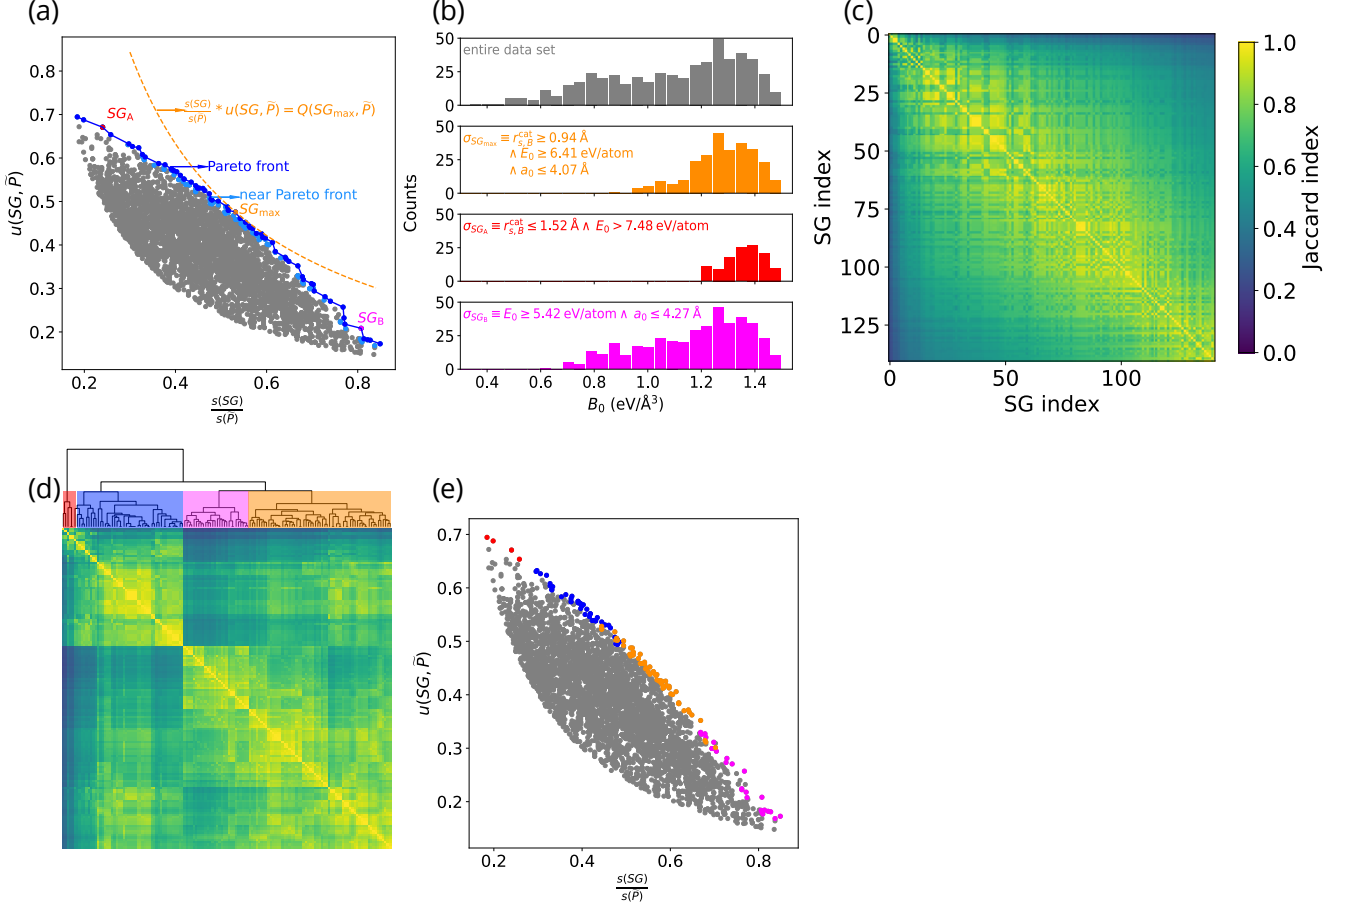
FIG. 1. Subgroup discovery of perovskites with high bulk modulus (B_0) using the positive mean shift as utility function $u(SG, \widetilde{P})$ (see Eq. 3) (a) Top 5,000 of the identified SGs (in grey), with the SG that maximizes Q highlighted in orange (SG_{max}) and the SGs belonging to the Pareto front (66 SGs) and Pareto region (141 SGs) highlighted in blue. Two examples of SGs at the Pareto front with lower and higher coverage compared to SG_{max} are highlighted in red and magenta, respectively (SG_{A} and SG_{B}). (b) From top to bottom: Distributions of the target B_0 in the entire data set, in SG_{max}, in SG_{A}, and in SG_{B}. The identified selectors for each SG are noted. They constrain the values of the radius of valence-s orbitals of $+1$ cations (cat) of B element ($r_{s,B}^{\mathrm{cat}}$), the equilibrium lattice constant (a_0) and the cohesive energy (E_0). (c) Similarity among the SGs belonging to the Pareto region. (d) Hierarchical cluster map obtained from the similarity matrix shown in (c). (e) Clusters of SGs obtained by the analysis of the dendrogram of (d) - see corresponding colors.

erties of the free atoms of the elements A and B as well as some bulk properties of the solid perovskite materials, namely the cohesive energy (E_0) and the equilibrium lattice constant (a_0). A data set of 504 materials, obtained using density functional theory calculations (DFT-PBEsol) is used (see ref. [15]). This considered material space includes A elements from the alkali, alkaline-earth, and scandium groups as well as lanthanides. The choice of B elements includes transition-metals as well as main-group elements such as bismuth, antimony, and germanium. Further details on the data set are provided in the Methods section.

Let us first analyze the results obtained with the

positive-mean-shift utility function, which is defined as

$$u(SG, \widetilde{P}) = \frac{\bar{y}(SG) - \bar{y}(\widetilde{P})}{y_{\mathrm{max}}(\widetilde{P}) - \bar{y}(\widetilde{P})}. \tag{3}$$

Here, $\bar{y}(SG)$ and $\bar{y}(\widetilde{P})$ are the mean values of the distribution of the target in the SG and in $\widetilde{P}$, and $y_{\mathrm{max}}(\widetilde{P})$ is the maximum value that the targets assumes in $\widetilde{P}$. This utility function requests that the values of the target within the SG are high with respect to $\bar{y}(\widetilde{P})$. It also assumes that the distributions of target values in the SG and in $\widetilde{P}$ are properly described by the mean values and that $\widetilde{P}$ is representative of the full population P.

The top-ranked SGs obtained with the utility function of Eq. 3 are shown as grey points in the utility function

vs. coverage plot of Fig. 1(a), along with the identified Pareto region in blue. The SGs belonging to the Pareto front and near it are highlighted in dark and light blue, respectively. This plot shows, in orange, the SG that maximizes Q, denoted $SG_{\max}$, and the curve corresponding to the constant value of $Q(SG_{\max}, \widetilde{P})$ as a dashed orange line. The Pareto region contains many SGs with Q values close to $Q(SG_{\max}, \widetilde{P})$ at intermediate coverage values. However, SGs at the Pareto region corresponding to lower and higher coverages such as SG_A and SG_B, shown in red and magenta respectively, present relatively lower Q values compared to $Q(SG_{\max}, \widetilde{P})$.

The analysis of the distributions of B_0 values in SG_A, $SG_{\max}$ and SG_B (Fig. 1(b)) indicates that, as the coverage increases, the SGs of the Pareto region have broader B_0 distributions and lower mean B_0 values. For the goal of identifying exceptional perovskites with high-B_0, the rules associated to SGs with low coverage and higher B_0 values are useful, since they provide a more focused description. Such SGs would not be detected based on the maximization of Q alone.

The selectors defining $SG_{\max}$, SG_A, and SG_B shown in the inset of Fig. 1(b), contain propositions constraining the values of the radius of valence-s orbitals of $+1$ cations (cat) of B element, $r_{s,B}^{\text{cat}}$. Moreover, these propositions constrain the value of a_0 and E_0 of the material to maximum and minimum thresholds, respectively. The propositions associated to $SG_{\max}$, SG_A, and SG_B are illustrated in the 2-dimensional plots of E_0 vs. a_0 and E_0 vs. $r_{s,B}^{\text{cat}}$ in Fig. 2. In these plots, the B_0 values for the materials in the dataset are indicated by the color of the circles. The identified inequalities describing SG_A, $SG_{\max}$, and SG_B highlight that materials with high E_0 and low a_0 present high B_0. This reflects the direct (inverse) relationship of B_0 with a_0 (E_0).[15]

In order to assess whether the identified SGs describe all relevant materials with high B_0 values, we verify whether the 26 perovskites presenting B_0 values higher than 1.43 eV/Å^3, the 95%-ile of the B_0 distribution over the training data set, are contained in $SG_{\max}$, SG_A, and SG_B. $SG_{\max}$ and SG_B, which are SGs with intermediate and high coverage values among the SGs in the Pareto region, respectively, contain all 26 materials with $B_0 \geq 1.43$ eV/Å^3. These materials are perovskites composed by the B elements chromium, manganese, iron, cobalt, and tungsten, and the A elements scandium, prseodymium, neodymium, cerium, promethium, yttrium, samarium, and beryllium. However, SG_A, which presents a lower coverage value but higher utility, does not include two of the 26 high-B_0 materials, those composed by the A element beryllium: BeMnO$_3$ and BeWO$_3$, with B_0 1.43 and 1.45 eV/Å^3, respectively. These two materials, highlighted in Fig. 2(b), present the lowest E_0 values (6.63 and 7.47 eV/atom, respectively) among the 26 materials. Thus, they do not satisfy the inequality $E_0 > 7.48$ eV/atom, which is part of the selector of SG_A (Fig. 1(b), in red).

The Jaccard similarity indices for all pairs of SGs of the Pareto region identified with the positive-mean-shift utility function are depicted by the colors in the similarity matrix of Fig. 1(c). In this figure, the SG indices are sorted according to increasing coverage. The regions close to the diagonal are associated to close-to-one $J(SG_i, SG_j)$, indicating that SGs with similar coverages values are similar. Moreover, this plot shows that SGs with low coverage present rather low similarity with the average- and high-coverage SGs. By applying the hierarchical clustering to the similarity matrix (Fig. 1(d)), distinct portions of the Pareto region are more clearly identified. In order to obtain examples of clusters that can identified in this analysis, we truncate the generated dendogram at the height indicated by upper edge of the colored rectangles in Fig. 1(d). We find four different clusters, colored in red, blue, magenta, and orange. These clusters are represented in the coverage vs. utility plot of Fig. 1(e) with the same colors as in Fig. 1(d). They roughly correspond to different coverage regimes. Thus, they present different levels of specificity. SG_A, $SG_{\max}$ and SG_B are examples of SGs belonging to the red, orange, and magenta clusters.

Finally, we evaluate the variability of the SGs with respect to the data set size by training the SGD with random selections of 75%, 50%, and 25% of the data set. Even though the Jaccard similarity indices decrease with decreasing data-set size, the SGs obtained with only 25% of the data set still present significant similarity compared to the SGs obtained with the entire data set. Thus, for the problem under consideration, SGD would be efficient even with a fraction of the data set. More details of this analysis can be found in the electronic supplementary information (ESI).

We now turn our attention to the utility function D_{cJS}, the divergence being evaluated between the distribution of the target values in the SG and the distribution of the target values in $\widetilde{P}$.[13] The information-theoretic D_{JS} is a symmetrized version of the Kullback-Leibler divergence (D_{KL}) or relative entropy. The D_{JS} between the discrete distributions R and S can be defined as

$$\begin{aligned} D_{\text{JS}}(R,S) &= \frac{1}{2}D_{KL}(R,M) + \frac{1}{2}D_{KL}(S,M) \\ &= \frac{1}{2}\sum_{x\in\chi} R(x)log\left(\frac{R(x)}{M(x)}\right) \\ &+ \frac{1}{2}\sum_{x\in\chi} S(x)log\left(\frac{S(x)}{M(x)}\right), \end{aligned} \quad (4)$$

where $M(x) = \frac{R(x)+S(x)}{2}$, and χ indicates the sample space. D_{JS} measures the dissimilarity between two distributions. It assumes small values for similar distributions and increases as the distributions are shifted with respect to each other or have different narrownesses. Thus, the D_{cJS} does not make assumptions on the shape of the distributions nor it explicitly requests that the values of the target in the SG are higher or lower compared
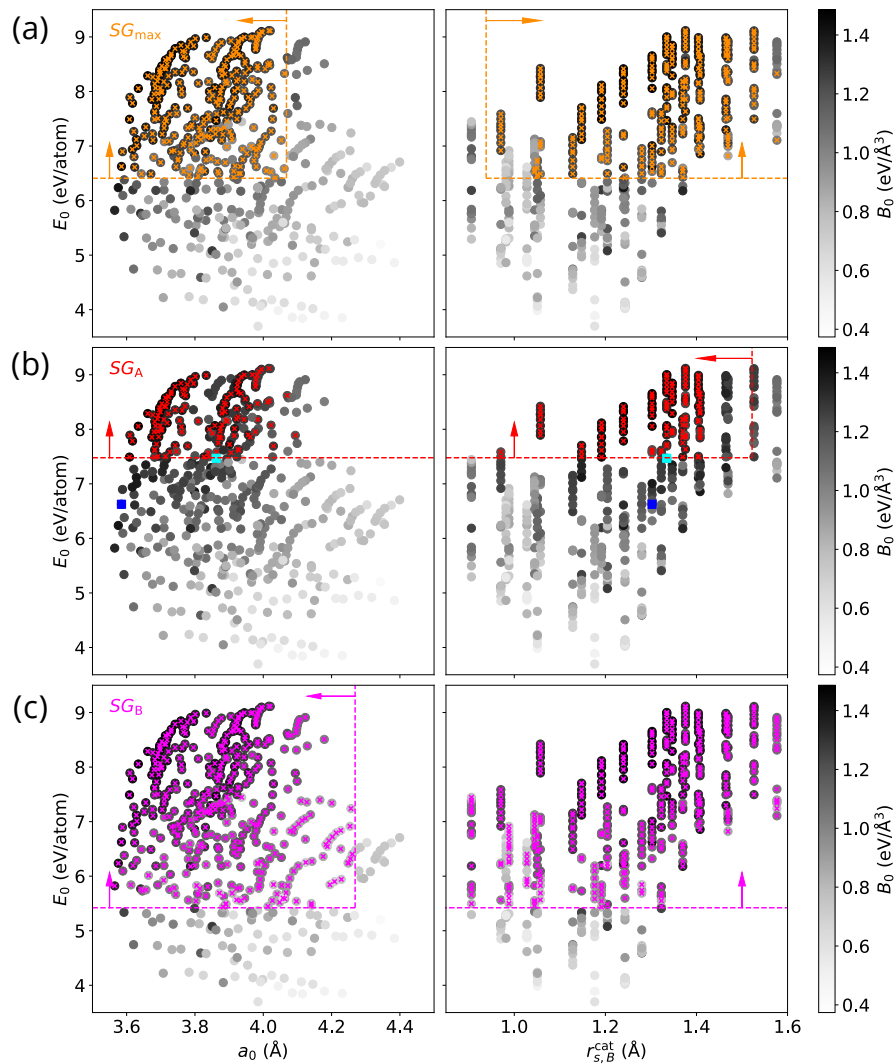
FIG. 2. Subgroup rules describing perovskites with high bulk modulus obtained using the positive mean shift utility function for (a) $SG_{\max}$, (b) SG_A, and (c) SG_B, as defined in Fig. 1(b). The propositions defining the SGs constrain the values of the equilibrium lattice constant (a_0), the cohesive energy (E_0) and the radius of valence-s orbitals of $+1$ cations (cat) of B element $r_{s,B}^{\mathrm{cat}}$, as indicated by the arrows. The grey scale of the circle colors indicates the value of the target property bulk modulus (B_0). The orange, red, and magenta crosses indicate the materials belonging to each of the identified SGs. The high-B_0 materials $BeMnO_3$ and $BeWnO_3$, which are part of the 95%-ile of the B_0 distribution over the entire data set but are not contained in the SG_A, are marked with blue and cyan squares in (b), respectively.

to the mean value in $\widetilde{P}$. Moreover, the D_{cJS} favors the selection of SGs presenting narrower distributions of target values than the distribution of the entire data set.

The results obtained with the $D_{cJS}(SG, \widetilde{P})$ utility function are shown in Fig. 3. The SGs identified at the Pareto region (Fig. 3(a)) contain materials with high B_0 (Fig. 3(b)). Moreover, these SGs present narrow distributions of target values, in particular the SGs at the low coverage region (e.g., SG'_A). Thus, by using the D_{cJS} utility function and the Pareto region concept, focused descriptions of high-B_0 materials can be identified in a systematic way. The selector defining SG'_A, shown in the inset of Fig. 3(b), contains propositions constraining the values of $r_{s,B}^{\mathrm{cat}}$, a_0, the expected oxidation state of

A in the perovskite (n_A), and the electron affinity and ionization potential of the element B (EA_B and IP_B, respectively). $SG'_{\max}$ and SG'_B are additional examples of SGs identified with the D_{cJS} utility function.

The 5% perovskites presenting the highest B_0 values of the data set are all contained in SG'_B. However, $SG'_{\max}$ does not contain the material $BeMnO_3$, which is part of the top 5% with respect to B_0. The SG'_A, in turn, does not contain three of the 26 high-B_0 materials: $BeMnO_3$, $BeWO_3$, and $ScWO_3$, with B_0 1.43, 1.45, and 1.44 eV/Å^3, respectively. The fact that SG'_A "misses" these materials can be related to the significantly different properties of beryllium and tungsten compared to the elements that are present in the remaining 23 high-
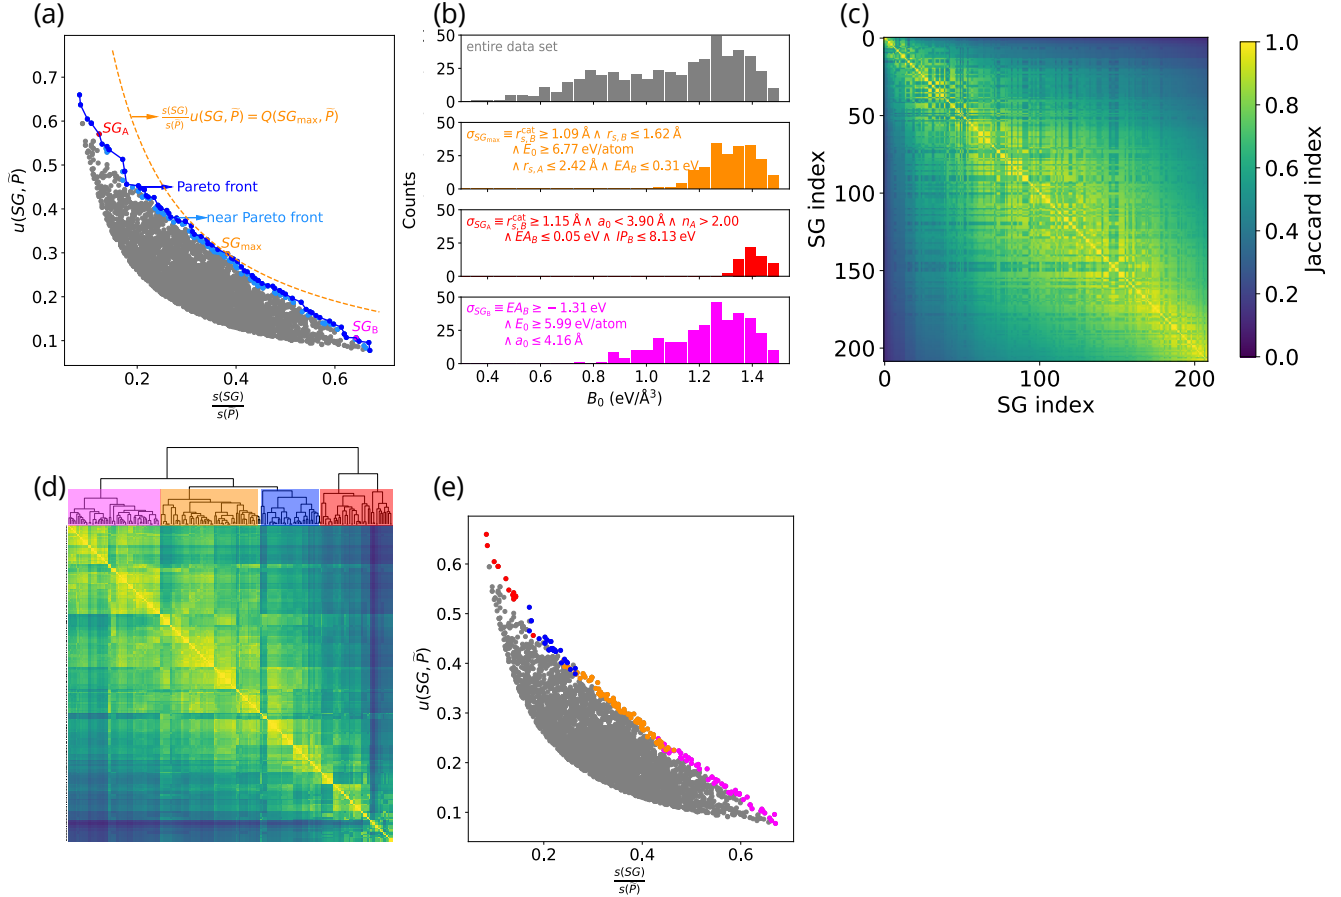
FIG. 3. Subgroup discovery of perovskites with high bulk modulus (B_0) using the cumulative Jensen-Shannon divergence (D_{cJS}) as utility function (see Eq. 4) (a) Top 5,000 of the identified SGs (in grey), with the SG that maximizes Q highlighted in orange ($SG'_{\max}$) and the SGs belonging to the Pareto front (79 SGs) and Pareto region (209 SGs) highlighted in blue. Two examples of SGs at the Pareto front with lower and higher coverage compare to $SG'_{\max}$ are highlighted in red and magenta, respectively (SG'_{A} and SG'_{B}). (b) From top to bottom: Distibutions of the target B_0 in the entire data set, in $SG'_{\max}$, in SG'_{A}, and in SG'_{B}. The identified selectors for each SG are noted. (c) Similarity among the SGs belonging to the Pareto region. (d) Hierarchical cluster map obtained from the similarity matrix shown in (c). (e) Clusters of SGs obtained by the analysis of the dendrogram of (d) - see corresponding colors.

B_0 materials, i.e., $3d$ metals (chromium, manganese, iron, cobalt) at the B site and metals from the scandium group or lanthanides at the A site.

The Jaccard similarity analysis and hierarchical clustering of the Pareto region identified with the D_{cJS} utility function (Fig. 3(c) and (d), respectively) highlight that the similarity score between the SGs with low coverage and the remaining SGs of the Pareto region is close to zero (e.g., the red cluster displayed in Fig. 3(d) and (e)). The Pareto front of SGs is sligthly modified when only a fraction of the data set is used for training the SG rules (see ESI for details).

Let us now compare the results obtained with the two utility functions. The Pareto region of SGs identified based on the positive mean shift utility function is associated to coverages in the approximate range [0.20, 0.80]

(Fig. 1(a)). The range of coverages of SGs in the Pareto region identified for the case of the D_{cJS} utility function is, in turn, ca. [0.10, 0.70] (Fig. 3(a)). Thus, optimal SGs with coverage values below 0.20, i.e., SGs that contain less than 20% of the dataset, could only be obtained with the the D_{cJS} utility function. These small SGs are also associated with narrower distributions of B_0 values. Thus, the rules corresponding those SGs are focused on the exceptional situations. This can be related to the fact that only the D_{cJS} explicitly favors narrow SGs.

We evaluated the similarity between SGs at the Pareto region obtained with the utilities positive mean shift and D_{cJS}. Figure 4(a) shows the similarity matrix obtained again using the Jaccard concept of Eq. 2. In this matrix, the SG indices are sorted from low to high coverages. The similarity between SGs with intermediate

coverages is high, as indicated by the yellow cells in Figure 4(a). However, SGs at low and high coverages are significantly different depending on the applied utility function. Indeed, the SGs with high coverages identified with the positive mean shift (e.g., SGs with indices ≥ 110 in Fig. 4(a), ordinate) and the SGs with low coverages identified with D_{cJS} utility function (e.g., SGs with indices ≤ 30 in Fig. 4(a), abssissa) do not correspond to similar SGs obtained with the other utility function. This is in line with the fact that the ranges of SG coverages achieved by each utility function are different (see above).

Finally, we compared in more detail the SGs that maximize the quality-function values according to the positive-mean-shift and Jensen-Shannon-divergence utilities, $SG_{\max}$ and $SG'_{\max}$, respectively. $SG_{\max}$ and $SG'_{\max}$, contain, respectively, 268 and 193 materials, as illustrated in Fig. 4(b) and (c). The Jaccard similarity index between these two SGs is 0.71 and there are 79 materials which are only part of one of these two SGs. In particular, $SG'_{\max}$ misses some high-B_0 materials, e.g., several ABO_3 perovskites composed by the B elements aluminum and titanium (in blue and green, respectively, in Fig. 4(c)), which have B_0 in the range 0.86-1.34 eV/Å^3 and are contained in $SG_{\max}$. Despite missing some of the high-B_0 materials, $SG'_{\max}$ presents overall a higher B_0 mean value than $SG_{\max}$(1.32 and 1.27 eV/Å^3, respectively) and a smaller narrowness (standard deviation of B_0 values equal to 0.09 and 0.12 eV/Å^3, respectively).

III. DISCUSSION

An approach for the identification of coherent collections of SGs based on a Pareto region was introduced. Using a measure for SG similarity, the SGs belonging to the Pareto region were hierarchically clustered. The sensitivity of SGD results with respect to the two utility functions positive mean shift and D_{cJS} was analyzed. The concepts were demonstrated by the learning of rules that describe perovskites with high bulk modulus. Our results show that rules focused on exceptional materials do not necessarily correspond to the SG that maximizes the quality function, but these rules can be identified with the Pareto region concept. The Pareto region analysis does not require additional computational effort, since the SGD solutions with quality-function values close to the maximum value are obtained on the fly during the optimization. We conclude that the D_{cJS} is a better choice of utility function than the positive mean shift for the purpose of identifying SG rules focusing on high-B_0 perovskites using the considered data set. Compared to the mean shift, the D_{cJS} has the advantage of simultaneously favoring SGs with shifted and narrow distributions of target values with respect to $\widetilde{P}$. The D_{cJS} utility function also alleviates the assumption that a summary value (e.g., the mean value) fully characterizes the distributions of target values in $\widetilde{P}$ and in the SGs, as the entire dis-

tributions are considered for evaluating D_{cJS}. Thus, the D_{cJS} utility function can handle distributions that deviate significantly from a Gaussian more efficiently than utility functions based on the mean shift. Nevertheless, D_{cJS} does not explicitly requests low or high values of target, which is a disadvantage compared to the positive (or negative) mean shift utility functions. The selection and testing of the utility function is thus crucial in SGD.

Finally, we note that the rules identified with SGD will be valid as long as the physical processes governing the materials in the training data set (and in the identified SGs) also govern the behavior of the materials in the materials space to be explored. In order to cover portions of the materials space where different underlying processes from those present in the training set are important, the incorporation of new data points and retraining of SG rules will be required.

The multi-objective perspective introduced in this contribution provides a framework for effectively dealing with the tradeoff between generality and exceptionality in SGD. This strategy combined with efficient algorithms for SG search and with a systematic incorporation of new data points to cover portions of materials space with significantly different characteristics than those of the training set will accelerate the AI-driven discovery of exceptional materials, while providing physical insights.

IV. METHODS

Data set As target materials property, we consider the bulk modulus (B_0), calculated using density-functional theory (DFT) with the PBEsol exchange-correlation functional. The candidate descriptive parameters used to characterize the cubic ABO_3 perovskites are shown in Table I. These features are atomic (elemental) properties of isolated atoms of the elements A or B such as orbital radii, ionization potential, electronegativity, etc., and properties of the solid perovskite materials equilibrium lattice constant (a_0) and cohesive energy (E_0). E_0 is defined as the energy needed per atom to atomize the crystal. In total, 11 candidate descriptive parameters were used. The detailed description of the DFT calculations is available in reference [15].

SGD approach We used the SGD approach as implemented in the realkd version 0.7.2. A Monte-Carlo-based SG search algorithm[5, 6] was used with 50,000 seeds for the initialization. The propositions were created based on 10 different thresholds per candidate descriptive parameter. These thresholds were determined by k-means clustering.

Pareto-region identification We consider the 5,000 SGs which are top-ranked with respect to Q. SGs are considered near the Pareto front if their distances to the Pareto front, measured in the coverage vs. utility plots, are lower than the threshold 0.01.

Hierarchical clustering An agglomerative hierarchical clustering is performed, using the SG similarity ma-
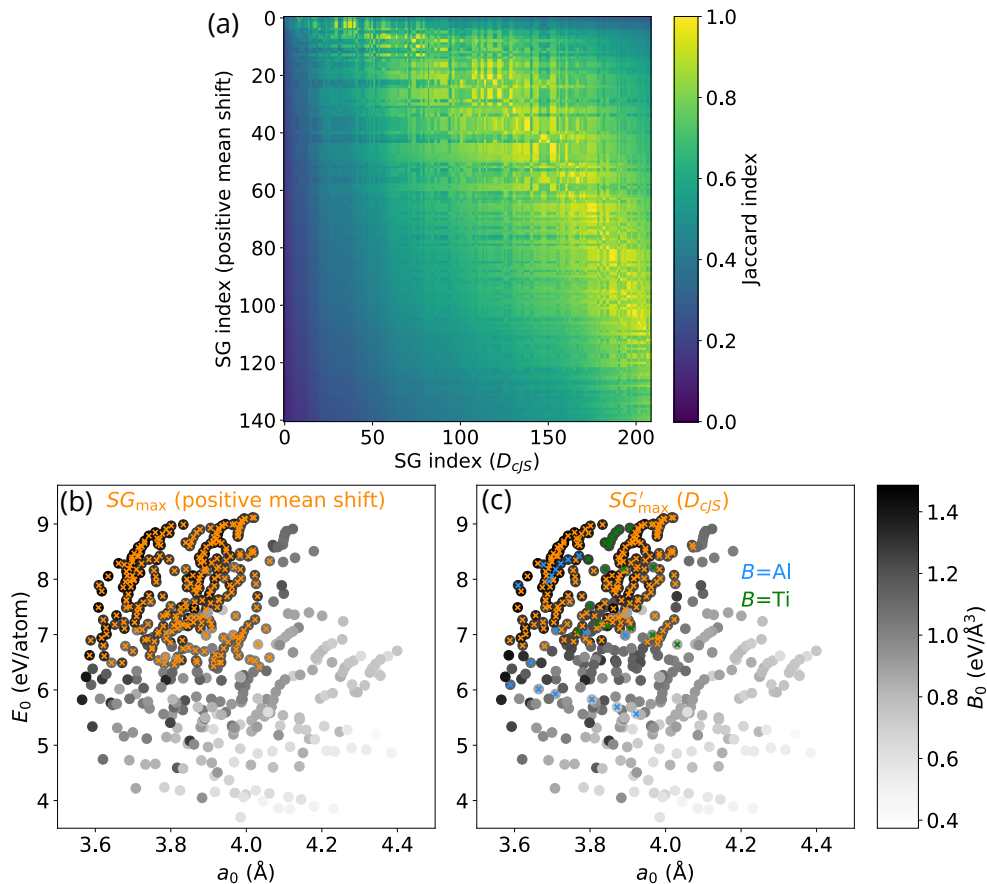
FIG. 4. (a) Jaccard similarity between SGs at the Pareto region obtained using the utility functions positive mean shift and cumulative Jensen-Shannon divergence (D_{cJS}). The SG indicies are sorted according to increasing coverage values. (b) and (c) illustrate the dataset in the cohesive energy (E_0) vs. equilibrium lattice constant (a_0) space along with the identified SGs that maximize the quality functions with the positive-mean-shift and cumulative-Jensen-Shannon-divergence utilities, SG_{max} and SG'_{max}, respectively. The grey scale of the circles indicates the value of the target property bulk modulus (B_0). The orange crosses indicate the materials in each of the identified SGs. The materials composed by the B elements aluminium and titanium, which are missed by SG'_{max}, are marked with blue and green crosses in (c), respectively.

TABLE I. Features or candidate descriptive parameters used to characterize the ABO_3 cubic perovskites. [a] Evaluated on an isolated atom using DFT-PBEsol. [b] Eveluted for the (solid) material using DFT-PBEsol.

Name	Symbol	Unit
Radii of the valence-s orbitals of the A and B atoms[a]	$r_{s,A}, r_{s,B}$	Å
Radii of the valence-s orbitals of the A and B +1 cations[a]	$r_{s,A}^{cat}, r_{s,B}^{cat}$	Å
Electron affinity of the A and B atoms[a]	EA_A, EA_B	eV
Ionization potential of the A and B atoms[a]	IP_A, IP_B	eV
Expected oxidation state of the element A	n_A	$\mathbb{Z}$
Equilibrium lattice constant[b]	a_0	Å
Cohesive energy[b]	E_0	eV/atom

trix defined by the Jaccard indices as input. The un-weighted pair group method with arithmetic mean was used as linkage criterion. The Euclidean distance was taken as distance measure. The hierachical clustering was performed using the seaborn package.[16]

V. DATA AVAILABILITY

The data set, including all input and output files of the calculations, are available in reference [15].

[1] Stefan Wrobel. An algorithm for multi-relational discovery of subgroups. In *European Conference on Principles of Data Mining and Knowledge Discovery*, 1997.

[2] Jerome H. Friedman and Nicholas I. Fischer. Bump hunting in high-dimensional data. *Stat. Comput.*, 9:123, Apr 1999.

[3] Bryan R Goldsmith, Mario Boley, Jilles Vreeken, Matthias Scheffler, and Luca M Ghiringhelli. Uncovering structure-property relationships of materials by subgroup discovery. *New Journal of Physics*, 19(1):013031, jan 2017.

[4] Mario Boley, Bryan R. Goldsmith, Luca M. Ghiringhelli, and Jilles Vreeken. Identifying consistent statements about numerical data with dispersion-corrected subgroup discovery. *Data Min. Knowl. Discov.*, 31(5):1391–1418, 2017.

[5] Mario Boley, Claudio Lucchese, Daniel Paurat, and Thomas Gärtner. Direct local pattern sampling by efficient two-step random procedures. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, page 582–590, New York, NY, USA, 2011. Association for Computing Machinery.

[6] Mario Boley, Sandy Moens, and Thomas Gärtner. Linear space direct pattern sampling using coupling from the past. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '12, page 69–77, New York, NY, USA, 2012. Association for Computing Machinery.

[7] Henrik Grosskreutz, Stefan Rüping, and Stefan Wrobel. Tight optimistic estimates for fast subgroup discovery. In Walter Daelemans, Bart Goethals, and Katharina Morik, editors, *Machine Learning and Knowledge Discovery in Databases*, pages 440–456, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[8] Lucas Foppa, Luca M. Ghiringhelli, Frank Girgsdies, Maike Hashagen, Pierre Kube, Michael Hävecker, Spencer J. Carey, Andrey Tarasov, Peter Kraus, Frank Rosowski, Robert Schlögl, Annette Trunschke, and Matthias Scheffler. Materials genes of heterogeneous catalysis from clean experiments and artificial intelligence. *MRS Bull.*, 46:1016–1026, Nov 2021.

[9] Christopher Sutton, Mario Boley, Luca M. Ghiringhelli, Matthias Rupp, Jilles Vreeken, and Matthias Scheffler. Identifying domains of applicability of machine learning models for materials science. *Nat. Commun.*, 11(1):4428, 2020.

[10] Lucas Foppa and Luca M. Ghiringhelli. Identifying outstanding transition-metal-alloy heterogeneous catalysts for the oxygen reduction and evolution reactions via subgroup discovery. *Top. Catal.*, 65:196, Jan 2022.

[11] Lucas Foppa, Christopher Sutton, Luca M. Ghiringhelli, Sandip De, Patricia Löser, Stephan A. Schunk, Ansgar Schäfer, and Matthias Scheffler. Learning design rules for selective oxidation catalysts from high-throughput experimentation and artificial intelligence. *ACS Catal.*, 12:2223, Jan 2022.

[12] Lucas Foppa and Matthias Scheffler. Towards a multi-objective optimization of subgroups for the discovery of materials with exceptional performance. 2023.

[13] Hoang-Vu Nguyen and Jilles Vreeken. Non-parametric jensen-shannon divergence. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2015, Porto, Portugal, September 7-11, 2015, Proceedings, Part II 15*, pages 173–189. Springer, 2015.

[14] Frank Nielsen. *Hierarchical Clustering*, pages 195–211. Springer International Publishing, Cham, 2016.

[15] Lucas Foppa, Thomas A. R. Purcell, Sergey V. Levchenko, Matthias Scheffler, and Luca M. Ghiringhelli. Hierarchical symbolic regression for identifying key physical parameters correlated with bulk properties of perovskites. *Phys. Rev. Lett.*, 129:055301, Jul 2022.

[16] Michael L. Waskom. seaborn: statistical data visualization. *Journal of Open Source Software*, 6(60):3021, 2021.