Matter

CellPress

Article

An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties

We introduce a general inverse design framework, called Fourier-transformed crystal properties (FTCP), which maps real- and reciprocal-space features onto properties using a VAE with property-structured latent space. We use FTCP to inversely design 142 unique crystals — not present in the training set or achieved by simple elemental substitution — with user-defined ground- and excited-state properties, including formation energy, bandgap, TE power factor, and combinations thereof. We validate our designed crystals using first-principles calculations, reaching a success rate of 7.1%–38.9%.

Zekun Ren, Siyu Isaac Parker Tian, Juhwan Noh, ..., Kedar Hippalgaonkar, Yousung Jung, Tonio Buonassisi

dannyzekunren@gmail.com (Z.R.)
buonassi@mit.edu (T.B.)

Highlights

We present an inverse-design framework for generative design of materials

The invertible crystal representation includes real- and reciprocal-space features

We generate 142 new crystals with target ground- and excited-state properties

We validate our results with first-principles calculations

2

Benchmark

First qualification/assessment of material properties and/or performance

Ren et al., Matter 5, 314–335
January 5, 2022 © 2021 Published by Elsevier Inc.
https://doi.org/10.1016/j.matt.2021.11.032

CellPress Matter

Article

An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties

Zekun Ren,[1,2,13,*] Siyu Isaac Parker Tian,[1,2,13] Juhwan Noh,[3] Felipe Oviedo,[4,14] Guangzong Xing,[5] Jiali Li,[6] Qiaohao Liang,[7] Ruiming Zhu,[8,9] Armin G. Aberle,[2] Shijing Sun,[4,16] Xiaonan Wang,[6,17] Yi Liu,[10] Qianxiao Li,[11] Senthilnath Jayavelu,[12] Kedar Hippalgaonkar,[8,9] Yousung Jung,[3] and Tonio Buonassisi[1,4,15,*]

SUMMARY

Realizing general inverse design could greatly accelerate the discovery of new materials with user-defined properties. However, state-of-the-art generative models tend to be limited to a specific composition or crystal structure. Herein, we present a framework capable of general inverse design (not limited to a given set of elements or crystal structures), featuring a generalized invertible representation that encodes crystals in both real and reciprocal space, and a property-structured latent space from a variational autoencoder (VAE). In three design cases, the framework generates 142 new crystals with user-defined formation energies, bandgap, thermoelectric (TE) power factor, and combinations thereof. These generated crystals, absent in the training database, are validated by first-principles calculations. The success rates (number of first-principles-validated target-satisfying crystals/number of designed crystals) ranges between 7.1% and 38.9%. These results represent a significant step toward property-driven general inverse design using generative models, although practical challenges remain when coupled with experimental synthesis.

INTRODUCTION

A common quest in materials research is to create a new material with a combination of user-specified properties, which is not present in any materials-property database. Historically, we may recruit an experienced scientist to use their intuition, to create a list of candidate compounds using heuristics (e.g., elemental substitution following a given set of rules). With the advent of materials-property databases, materials screening became commonplace. Theoretical screening of solid-state materials using elemental substitution and mixing has allowed for the discovery of several crystals with user-defined functional properties, e.g., perovskite materials with tailored bandgaps.[3–8] However, even under the high-performance computing (HPC) framework, the computational cost of density functional theory (DFT) calculations is high, prohibiting an exhaustive search of the theoretical materials space.[9,10] Consequently, the leading databases contain on the order of 10^5–10^6 calculations for solid materials[2,11,12] — only a tiny fraction of the number of stoichiometric inorganic compounds believed to be possible (order 10^{10} considering quaternary crystals).[13,14]

To overcome these limitations, machine-learning (ML) methods have been developed to inversely design crystalline solids. (Inverse design refers to the act of a

Progress and potential

Inverse design refers to designing a new material with user-specified properties, while general inverse design, in addition, entails the designed materials to be of various chemistries and structures. Hitherto, the proposed inorganic generative-model-based inverse design algorithms are limited to generate crystals within a given chemical system, e.g., V_xO_y, or a given crystal structure, e.g., cubic structures. Realizing the general aspect of inverse design requires an invertible representation that embeds both chemistry and structure (generalized), as demonstrated herein. The general inverse design framework proposed here features a generalized invertible representation, accessing both chemistry and structure, and a property-structured latent space from a VAE, enabling the generative capacity. The proof-of-concept demonstration of this framework in designing for user-specified ground- and excited-state properties marks a step toward general inverse design.

user defining target material properties and inferring a material that meets target properties, e.g., by using an algorithm).[15,16] For the ML model, there are two major approaches to inversely design crystalline solids used today: global optimization and generative models. Global optimization, also called "directed evolution," involves modifying known or randomly enumerated compounds using a set of rules to design new compounds; their exploratory capacity is limited by the initial selection of structures and elements. Generative models, which learn a given data distribution, directly model all training-set materials into a probabilistic representation, from which new materials can be sampled.

Two most commonly used generative models in inverse design of solid-state materials are generative adversarial network (GAN)[17] and variational autoencoder (VAE),[18] and the key enabler of both, on top of the algorithm itself, is an invertible crystallographic representation. By "invertible", we refer to the both-way conversion from materials to representation, and vice versa. Especially, the conversion from representation to materials (the inverse of the one-way conversion from material to representation, enabling property prediction) requires a materials representation that enables the algorithm to extract necessary crystallographic information, for instance, the site location and occupancy information contained in a crystallographic information file (CIF), an input file format of choice for DFT calculations.

Because of the difficulty in creating a general invertible crystallographic representation, early demonstrations of inverse design using generative models were often limited to a fixed subset of elements[7,8,14,19,20] or crystal structures.[21–23] We observe that there lacks a general inverse design framework using generative models for inorganic crystals in prior art. Our definition of "general inverse design" is the ability to produce a prediction of a specific material (both chemistry and structure) on the basis of a user-specified target property (or properties), i.e., solving the inverse problem of property prediction. To perform inverse design using ML, as opposed to human intuition, we posit that a materials representation is required that must include two key elements: (1) a representation that incorporates both structure and chemical composition into the descriptor (both structure and composition varying), and (2) a representation that is invertible, making it amenable to solving the inverse problem (property → structure + chemistry). Current invertible representations (Figure 1, and Table S1) are representative of either only specific compositional spaces (not composition varying), such as V_xO_y space,[7,8,14,19,20] or only specific structural spaces (not structure varying), such as cubic structures.[21–23] Korolev et al. proposed a spectrum representation of composition + powder X-ray diffraction (XRD) pattern;[24] although allowing variation of both composition and structure, the conversion from representation to material, i.e., the construction of a unit cell from the composition + powder XRD pattern, is hard to achieve with algorithmic automation, but may still be possible for experienced human experts (thus our rating of "Limited/No" for its invertibility in Table S1). An outstanding challenge in the field is to develop an invertible crystallographic representation that accesses various chemical systems (composition varying) and various crystallographic space groups (structure varying), thus enabling general (and property-driven) inverse design.

In this work, we propose a framework for general inverse design of inorganic crystals, which is orders-of-magnitude faster than naive screening of the entire chemical space. The framework features two main components: (1) A generalized crystallographic representation, which is both composition and structure varying. The representation comprises (A) real-space features, containing CIF-like information, thus

[1] Low Energy Electronic Systems (LEES), Singapore-MIT Alliance for Research and Technology (SMART), Singapore 138602, Singapore

[2] Solar Energy Research Institute of Singapore (SERIS), National University of Singapore, Singapore 117574, Singapore

[3] Department of Chemical and Biomolecular Engineering (BK21 Four), Korea Advanced Institute of Science and Technology (KAIST), 291 Daehakro, Daejeon 34141, Korea

[4] Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139-4307, USA

[5] Okinawa Institute of Science and Technology Graduate University (OIST), Okinawa 904-0495, Japan

[6] Department of Chemical and Biomolecular Engineering, National University of Singapore, Singapore 117585, Singapore

[7] Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[8] Institute of Materials Research and Engineering, Agency for Science, Technology and Research (A*STAR), Singapore 138634, Singapore

[9] Department of Materials Science and Engineering, Nanyang Technological University, Singapore 117575, Singapore

[10] Materials Genome Institute (MGI), Shanghai University, Shanghai 200444, China

[11] Department of Mathematics, National University of Singapore, Singapore 117543, Singapore

[12] Institute for Infocomm Research, Agency for Science, Technology and Research (A*STAR), Singapore 138632, Singapore

[13] These authors contributed equally

[14] Present address: Microsoft AI for Good, Redmond, WA 98052, USA

[15] Lead contact

[16] Present address: Toyota Research Institute, Los Altos, CA 94022, USA

[17] Present address: Department of Chemical Engineering, Tsinghua University, Beijing 100084, China

*Correspondence:
dannyzekunren@gmail.com (Z.R.),
buonassi@mit.edu (T.B.)

https://doi.org/10.1016/j.matt.2021.11.032

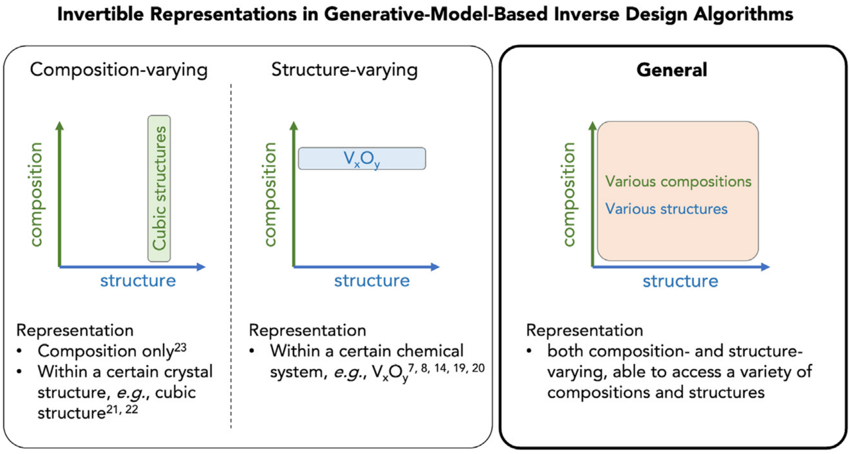

Figure 1. Novelty of the FTCP inverse design framework (right), compared to prior art (left)
The FTCP framework allows for both composition and structure to vary, enabling general and property-driven inverse design. See also Table S1.

guaranteeing the invertibility, and (B) reciprocal-space features, a reciprocal-space formulation of crystal properties similar to the structure-factor (Fourier transform) calculation in XRD, as an additional featurizer. The latter is inspired by Ziletti et al.'s representation of 2D diffraction fingerprint,[25] which leverages the more compact crystal periodicity and symmetries in the reciprocal space. (2) A VAE, with an extra target-learning branch connected to the latent space, mapping latent vectors/points to user-specified property(ies); the encoder encodes the crystals in the training set into a continuous probabilistic latent space, and the decoder decodes any vector in the latent space to its corresponding crystals (sampling the latent space for designing new crystals); the target-learning branch during training will jointly organize the latent space to reflect a continuous change in property, i.e., a property gradient, and we term thus-formed latent space "property structured." We name both the generalized invertible representation (1) and the overall framework (1)+(2) after the reciprocal-space featurizer, as the Fourier-transformed crystal properties (FTCP) framework (sometimes simply referred to as FTCP).

Using FTCP, we demonstrate three inverse design cases with design targets from single to multiple, from simple to complex: (1) case 1, designing for various targeted formation energies, E_f (ranging from -0.3 to -0.7 eV/atom); (2) case 2, designing for bandgap, $E_g = 1.5$ eV, while keeping $E_f < -1.5$ eV/atom (a desirable bandgap target for photovoltaic and optoelectronic applications); (3) case 3, designing for thermoelectric (TE) power factor (an excited-state property) to be as large as possible, while keeping $0.3 < E_g < 1.5$ eV and $E_f < 0$ eV/atom (a desirable power factor target for high-efficiency TE materials and a desirable bandgap range for a low- and medium-temperature range). The designed crystals are unique, i.e., not in the Materials Project[2] database (from where our training and test sets derive), and they span a variety of chemistries and crystal structures (i.e., both composition and structure varying). We "validate" the designed crystals using DFT structural relaxation followed by property prediction using first-principles calculations and define a designed crystal as "successful" if it relaxes properly and achieves the user-specified target property to within a given range. To compare FTCP against a baseline, we calculate the "random success rate" as the probability of finding a material with the user-specified target property by randomly picking from the training set, and we quantify the improvement of FTCP over random.

A summary of main results presented in the paper follows: (1) case 1, FTCP achieves success rates (number of target-satisfying designed crystals/number of total designed crystals) from 14.3% to 38.9%, scoring improvement over random from 38.8% to 270%; (2) case 2, success rate is 36.8%, achieving improvement over random of 560%; (3) case 3, FTCP designs two unique crystals that achieve comparable peak power factors with the state-of-the-art TE material, germanium telluride (GeTe), scoring a success rate of 7.1% (improvement over random not quantified due to the lack of complete power factor labels).

With the above cases, we demonstrate the usefulness of FTCP and validate using first-principles calculations, as is common practice in prior art (references are detailed in the section synthesizability challenge). We posit that a complete experimental validation is outside of scope for this study, but to lay the groundwork toward the experimental-synthesis goal of inverse design, we explore adding a näive synthesizability metric, i.e., the existence or absence of an Inorganic Crystal Structure Database (ICSD) entry (previously explored by Jang et al.),[26,27] to the target-learning branch. This organizes the property-structured latent space according to a synthesizability metric, which can direct generative-design sampling toward more favorable regions (of higher likelihood to possess an ICSD entry). We conclude with a brief discussion of the invariance challenge faced by structure-conscious invertible crystallographic representations, including the FTCP representation.

RESULTS AND DISCUSSION

Inverse design framework, FTCP

Representation

Representing an infinitely repeating 3D crystal is hard, compared with representing small organic molecules, due to the periodicity, the complex geometrical symmetries (230 space groups for 3D periodic crystals), and the vast chemical space (up to 81 nonradioactive elements). To sufficiently capture this variety of crystals while satisfying invertibility, we propose a crystallographic representation of real-space CIF-like features combined with reciprocal-space Fourier-transformed features (Figure 2A).

The real-space features contain necessary information for unit cell construction in 2D matrix form, namely a vertical concatenation of:

- Element matrix, where each column is a one-hot vector representing a constituent element.
- Lattice matrix, where the lattice parameters, a, b, c, α, β, and γ, form a 2×3 matrix.
- Site coordinate matrix, where each row vector contains the fractional coordinate of a site (vacant sites have all-zero entries, namely zero padding).
- Site occupancy matrix, or site index matrix, where each row is a one-hot vector for a site indicating which constituent element occupies the site (vacant sites have all-zero entries).
- Elemental property matrix, where each column is an elemental property vector/elemental descriptor $\boldsymbol{Z}$ (equivalent to the atom feature vector used in the crystal graph convolutional neural network (CGCNN)[28] representation) of the constituent element, including group number, period number, electronegativity, covalent radius, and valence electrons, among others.

The real-space features guarantee invertibility because if the real-space features are set, i.e., filled with values, a crystal can always be constructed. The resultant crystal

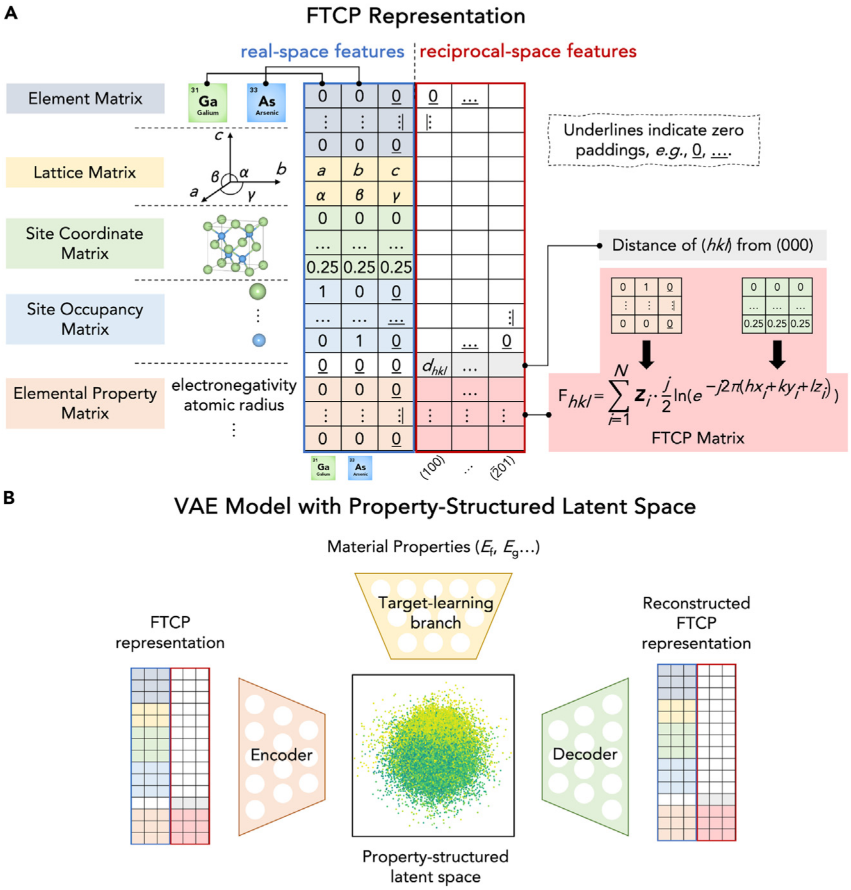

Figure 2. FTCP framework, representation + model

(A) Representation: The invertible FTCP representation contains both real- and reciprocal-space features. Real-space contain CIF-like features, such as the element matrix (describing constituent elements), the lattice matrix (describing lattice parameters), the site coordinate matrix (describing fractional coordinates of sites), the site occupancy matrix (describing elemental occupancy per site), as well as the elemental property matrix (elemental descriptors). Reciprocal-space features project the elemental descriptor $\boldsymbol{Z}_i$ (i for each site) for all the N sites in the unit cell along various spatial frequencies, hkl (Miller indices) via spatial discrete Fourier transform (with a data preprocessing of $\frac{j}{2}$ln where j is the imaginary unit) to form the FTCP matrix. The distance of each hkl (k point) from (000) (light-gray row, with "d_{hkl}", in reciprocal-space features) is also recorded and prepended to the FTCP matrix (red box in reciprocal-space features). (GaAs unit cell is used as an example for illustration).

(B) Model: the VAE architecture using the invertible FTCP representation for inverse design. On top of the encoder + decoder architecture of a normal VAE, the latent space is also connected to a target-learning branch for property mapping, reflecting a property gradient(s) (property-structured latent space). The latent space is visualized by plotting two dimensions (the third and the 12th) out of a total of 256. A detailed discussion of the visualization can be found in section S2.1 in supplemental information.

may not valid, but there is always a corresponding resultant crystal. Analogous to a CIF, the values in the real-space representation always correspond to a specific crystal, although the crystal itself may not be valid.

The reciprocal-space (also called k-, momentum-, frequency-, and Fourier-space) features contain the spatial Fourier-transformed elemental property vectors (the transformed property is termed crystal property). The formulation of the reciprocal-space features roots in the domain knowledge of solid-state physics. Atoms in a crystal are arranged in a periodic pattern. Bloch's theorem establishes that the wave function can be expressed as the product of a plane wave and a function that has the same periodicity as the crystal.[29] The spatial periodicity of the atomic arrangement, analogous to the temporal periodicity of a signal, can be analyzed in the frequency domain with Fourier transform, or in materials science terms, the reciprocal space. We build upon the discrete spatial Fourier transform used in the structure factor calculation in scattering physics and XRD to reach Equation 1:

1. Instead of transforming the atomic scatter factor to the reciprocal space, we substitute with the elemental property vector (projecting elemental properties to obtain crystal properties).
2. Apply a preprocessing ($\frac{j}{2}\ln$ increases the prediction accuracy for the formulations attempted in this study; see details in section S2.2 in supplemental information).

$$F_{hkl} = \sum_{i=1}^{N} Z_i \bullet \frac{j}{2} \ln\left(e^{-j2\pi(hx_i+ky_i+lz_i)}\right) = \sum_{i=1}^{N} Z_i \bullet \pi(hx_i+ky_i+lz_i) \qquad \text{(Equation 1)}$$

where hkl are the Miller indices/spatial frequencies, Z_i is the elemental property vector for site i (total N sites in the unit cell), $x_i y_i z_i$ are the fractional coordinates of site i, and j is the imaginary unit. The FTCP, F_{hkl}, is calculated for 59 different combinations of hkl (k points), which, as column vectors, concatenate horizontally to form an FTCP matrix. Upon prepending the distance of the k points from (0, 0, 0) to the FTCP matrix, the reciprocal-space features (matrix) become complete. Horizontal concatenation of the real-space features (matrix) and the reciprocal-space ones (with zero padding) forms the FTCP crystallographic representation. (The reciprocal-space features act as an additional featurizer, which improves the reconstruction accuracy of the VAE + target-learning branch model. See details in section S2.3 in supplemental information.)

Model

After obtaining the invertible FTCP representation (2D matrix), we construct a continuous property-structured latent space, which allows for sampling of new crystals using a VAE. The VAE comprises an encoder, a decoder, and a target-learning branch, as shown in Figure 2B. The encoder compresses the FTCP representation into a point probabilistically in the latent space of reduced dimension (with total dimensions = 256), namely the encoder outputs z_{mean} and z_{variance}. The decoder learns to sample with z_{variance} around the latent point z_{mean} and to decode this vicinity of z_{mean} back to the original FTCP representation. Naturally, the training for better reconstruction (minimizing reconstruction loss) results in nearness of latent points of similar FTCP representations, namely a cluster-forming behavior in the latent space for similar inputs. We also implement a standard VAE Kullback-Leibler (KL) loss to promote dense-packing of latent points around the center of the latent space, especially for those of very different FTCP representations, by encouraging z_{mean} and z_{variance} to follow that of a unit Gaussian. The reconstruction loss and the KL loss forms a stand VAE, where the input is embedded to a continuous probabilistic latent space — the cluster forming of the probabilistic points results in continuous change in the decoded FTCP representation, and the dense packing ensures the latent space does not contain void areas without encoded latent points. Overall, the two losses ensure a continuous latent space, which allows for easy sampling between known crystals (encoded latent points of training data). In addition, we add a feedforward target-learning branch to map latent points to certain properties, namely there is an additional property-mapping loss. (The target-learning branch handles multiple target properties by outputting a vector, instead of a scalar, containing entries for each property. Thus, the property-mapping loss is the overall loss across different properties.) This additional loss further encourages latent points to be near when they possess similar property values, thus introducing a property gradient (change) in the latent space (organizing latent points further according to the property). We term the obtained latent space "property structured." The overall loss of the VAE is a weighted sum of the three losses, i.e., the reconstruction loss, the KL loss, and the property-mapping loss. The detailed architecture and the hyperparameters of the model are recorded in the section generative model and section S2.4.

Design from trained model (sampling and postprocessing)

To sample from the trained property-structured latent space, we adopt a local perturbation (Lp) strategy, similar to that used by the decoder during training. We identify the latent points of target-satisfying crystals in the training set, and we sample around the latent points with a scaled (0.3–3) unit Gaussian noise, i.e., a Gaussian noise with scaled variance. The scale controls a tradeoff between exploitation and exploration: with too small a scale, the sampled points fall too close to known crystals, displaying high reconstruction accuracy with low novelty; with too large a scale, the sampled points are more scattered, displaying high novelty with low reconstruction accuracy. Thus, our choice of the scale is determined through trial and error to maximize exploration, while ensuring sufficient exploitation by taking the largest value that yields structure design errors <20%, discussed in greater details in section S2.5 in supplemental information. Apart from Lp, we also explore two other sampling strategies, spherical linear interpolation[30] (Slerp) and global perturbation (Gp), in section S2.5 in supplemental information. We deem Lp to be the most suitable and high-achieving strategy in the setting of property-driven sampling, and we consistently use Lp in the design (use) cases in this study.

The decoding from sampled latent points results in a decoded/reconstructed FTCP representation (real-space + reciprocal-space features). The reconstruction of a unit cell lies in postprocessing the decoded real-space features, which entails the processing of the element matrix (one-hot encoded), the lattice matrix, the site coordinate matrix, and the site occupancy matrix (one-hot encoded). For one-hot encoded matrices, the respective one-hot vector is obtained by setting the maximum value in the vector to one and the rest to zero. For the site occupancy matrix, to distinguish empty sites due to zero padding, we enforce a threshold (0.05) such that empty sites are where all decoded values in the vector are below the threshold. Thus, empty sites will have all-zero vectors in the site occupancy matrix, while valid sites will have one-hot vectors after postprocessing. We then use the site occupancy matrix to ignore the decoded coordinates for empty sites.

Workflow

Our workflow for the design cases (applying FTCP) in this study is shown in Figure 3, where we have four stages: (1) We define the target property of the intended material, (2) we design from trained model (performing sampling and postprocessing) and obtain a number of FTCP-designed candidates, (3) we pass these candidates through structural relaxation (performed by DFT in our study; in the future, this may be performed by an ML surrogate model to approximate the speed of FTCP compound generation), and we remove the nonconverged, repeated, and invalid structures, e.g., with overlapping atoms. (DFT structural relaxation is a way to reduce

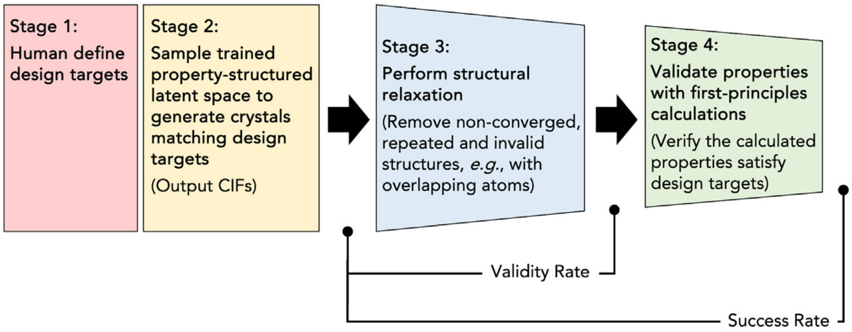

Figure 3. Workflow of the FTCP framework in design cases

Note that structural relaxation can be performed using DFT (our study), or in the future, an ML surrogate model. We define the following metrics: validity rate, percentage of FTCP-designed candidates successfully passing through structural relaxation ([no. of candidates exiting stage 3]/[no. of candidates exiting stage 2]), and success rate, percentage of FTCP-designed candidates confirmed to have the user-specified target property(ies) to within the specified margin of error ([no. of candidates exiting stage 4]/[no. of candidates exiting stage 2]).

errors in the construction of designed crystals. We discuss the sources of error and the need for structural relaxation for designed crystals in section S2.6 in supplemental information.) (4) We perform first-principles calculation(s) to verify the property(ies) of the designed candidates and retain those that satisfy the design target within a user-specified margin of error (tolerance).

We then define three metrics to quantify the performance of FTCP: (1) Validity rate, defined to be the percentage of FTCP-designed candidates successfully passing through the DFT structural relaxation ([no. of candidates exiting stage 3]/[no. of candidates exiting stage 2]). (2) Success rate, defined to be the percentage of FTCP-designed candidates confirmed to have the user-specified target property within tolerance ([no. of candidates exiting stage 4]/[no. of candidates exiting stage 2]). (3) Improvement over random, where random success rate is the probability of finding a material with the user-specified target property by randomly picking it from the training + test datasets [(FTCP success rate – random success rate)/random success rate]. We approximate the random success rate with the percentage of target-satisfying crystals in the training + test datasets.

With the above workflow, we apply FTCP in three design cases:

1. Design for formation energy, E_f: We design based on ternary crystals for four E_f targets, −0.5, −0.3, −0.6, and −0.7 eV/atom, as these E_f are the most prevalent in the database we use for training (Materials Project[2]). To calculate the "success metric," we define a tolerance of ±0.06 eV/atom.
2. Design for bandgap, E_g, with E_f constraint: We design based on ternary and quaternary crystals for a design target E_g = 1.5 eV, which is of interest both for solar cells (near the maximum of the detailed balance efficiency limit for single-junction devices[31]), and for optoelectronic devices (e.g., LEDs). We also set a negative E_f (<-1.5 eV/atom) target as a proxy for the stability of designed crystals. To calculate the "success metric," we define a tolerance of ±0.3 eV for E_g and ±0.06 eV/atom for E_f.
3. Design for TE power factor, with E_g and E_f constraints: We design based on ternary and quaternary crystals for a power factor (heat-to-electricity conversion efficiency) to be as large as possible. We also set a bandgap target between 0.3 and 1.5 eV, desirable for low- and medium-temperature range, as

Table 1. Performance of three design cases using FTCP

	Case 1				Case 2	Case 3
Description	Formation energy (±0.06 eV/atom tolerance)				Bandgap (±0.3 eV tolerance) with formation energy constraint	TE power factor (as high as possible) with bandgap and formation energy constraints
Metrics	$E_f = -0.5$ eV/atom	$E_f = -0.3$ eV/atom	$E_f = -0.6$ eV/atom	$E_f = -0.7$ eV/atom	$E_g = 1.5$ eV $E_f < -1.5$ eV/atom	Power factor as large as possible 0.3 eV $\leq E_g \leq 1.5$ eV $E_f < 0$ eV/atom
Validity rate	77.8% (14/18)	81.0% (17/21)	96.4% (27/28)	92.9% (26/28)	84.2% (16/19)	42.9% (12/28)
Success rate (FTCP)	38.9% (7/18)	14.3% (3/21)	17.9% (5/28)	21.4% (6/28)	36.8% (7/19)	7.1% (2/28)
Success rate (Random)[a,b]	10.5% (2,781/26,402)	10.3% (2,732/26,402)	9.6% (2,522/26,402)	8.3% (2,183/26,402)	5.5% (3,035/54,925)	–[c]
Improvement over random[b]	270%	38.8%	86%	150%	560%	–[c]

[a] approximated by the percentage of target-satisfying crystals in the dataset (training + test).

[b] calculated with an updated version of Materials Project accessed on 14 Sep 2021. (Materials Project has updated since our design cases, which accessed Materials Project on 22 Jun 2020.)

[c] not calculated for case 3 because of the lack of calculated power factor values for every crystal and the use of a qualitative criterion, "power factor as large as possible".

well as a negative (<0 eV/atom) formation energy, again as a proxy for stability. To calculate the "success metric," we compare the first-principles calculated power factor to the state of the art TE material, GeTe, and define a tolerance of ±0.3 eV for E_g and ±0.06 eV/atom for E_f.

After obtaining the designed candidate CIFs, according to our workflow, we pass the designed candidates through DFT structural relaxation to remove nonconverged, repeated, and invalid structures (validity rate is calculated). We then apply DFT calculations for E_f, and E_g, and Boltzmann transport properties (BoltzTraP[32]) calculations for power factor. (The details of the structural relaxation, and the first-principles calculation of properties, are described in sections S3.2 and S3.3, and we provide a list of validated FTCP-designed crystals in S3.1.) We verify the designed candidates to retain those with target-satisfying properties within aforementioned tolerances (success rate is calculated).

We summarize the performance of the design cases in Table 1 using the metrics defined above (detailed descriptions of each design case are in the following section applying FTCP: three design cases validated by first-principles calculations).

We observe a rule of thumb that with more design targets (increased design difficulty), the percentage of target-satisfying crystals in the dataset (random success rate) decreases, and the validity rate and the success rate also both decrease because there are more losses (property-mapping losses) to optimize. As the inverse design problem becomes more complex (e.g., the number of target properties increases, the tolerance tightens), the usefulness of FTCP is expected to grow; see the improvement over random for case 2 versus case 1, for example.

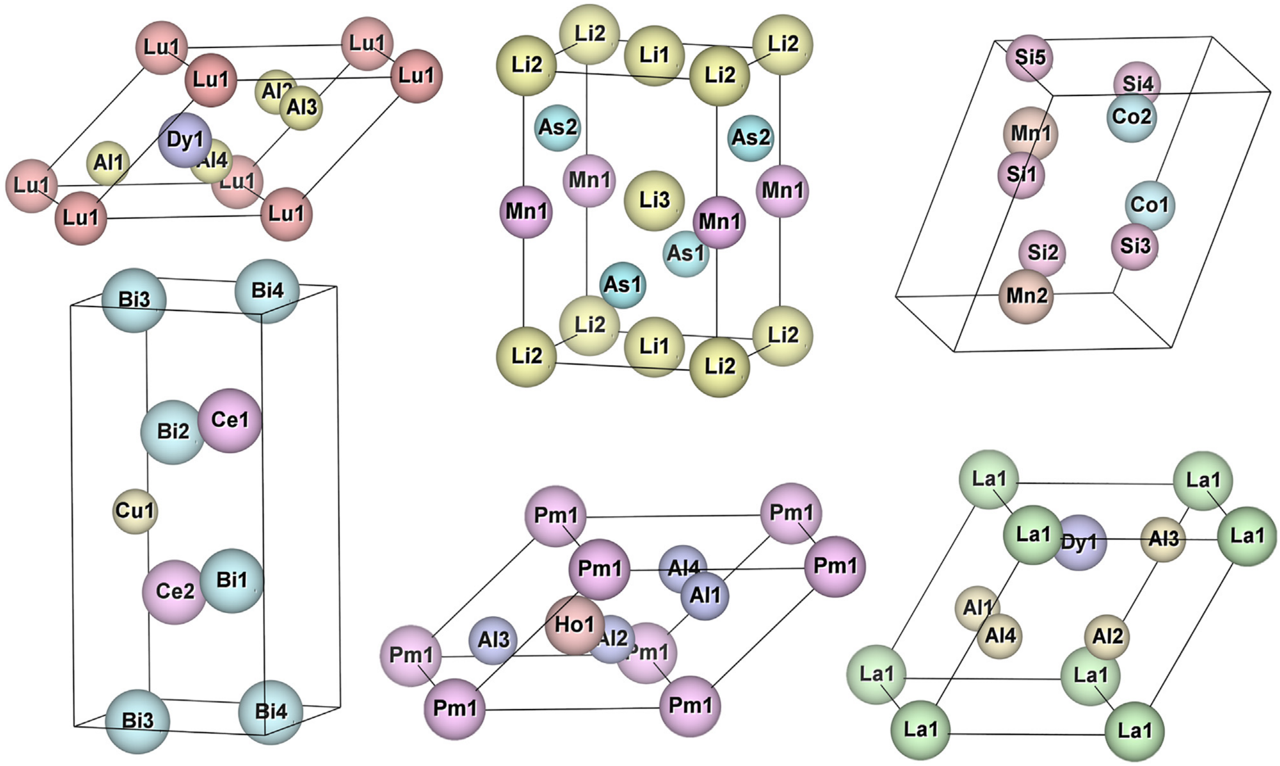

Figure 4. Examples of FTCP-designed crystals
The designed crystals are with targeted $E_f = -0.5$ eV/atom after DFT relaxation. See also Figure S4.

There are two sources of error reducing success rate: one is from imperfect positioning of latent points and imperfect reconstruction (KL loss and reconstruction loss), and one is from imperfect property mapping (property-mapping loss), which can yield three scenarios of error: (1) The validity of the decoded crystal is too low that the geometry is wrong or the structural relaxation fails. (2) The decoded crystal successfully relaxes, but the resultant crystal is not what is intended after relaxation (thus, the first-principles calculation of the resultant crystal yields the incorrect property value). (3) The property mapping is not accurate, resulting in an inaccurate property gradient (thus, the first-principles calculation yields an incorrect property value). Scenarios 2 and 3 can also happen at the same time. We posit that although a high success rate is ideal, the main achievement of FTCP is realizing a nonzero success rate of general inverse design. Many research problems only require one or a few successful candidate(s); the minimum success rate is therefore on the order of the inverse of the throughput of the synthesis tool.

Applying FTCP: Three design cases validated by first-principles calculations

Case 1: Design for formation energy and case 2: Design for bandgap (with formation-energy constraint): Case studies for photovoltaic and optoelectronic applications

In the first design case (case 1) of the FTCP framework, we design new crystals based on ternary crystals with target $E_f = -0.5$ eV/atom. As our dataset (training + test), we select the ternary crystals of ≤ 20 sites with energy above hull <0.08 eV/atom in the Materials Project database. Examples of six FTCP-designed 3D crystals with targeted E_f after DFT relaxation are shown in Figure 4. These crystals cannot be found in the Materials Project database, indicating their uniqueness. We perform DFT calculations of structural relaxation and formation energy using GGA(+U). The

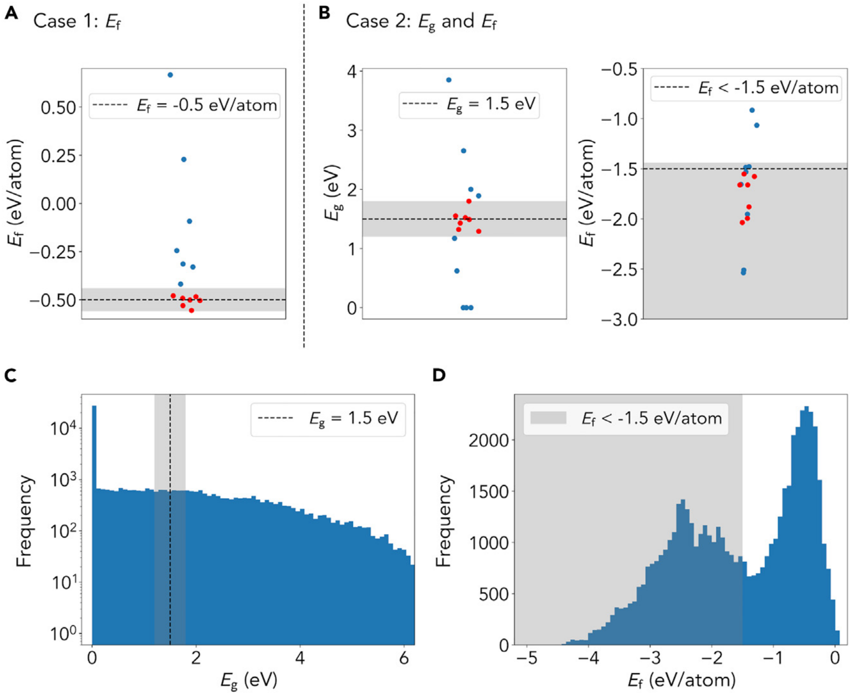

Figure 5. DFT-calculated properties of FTCP-designed crystals and target-satisfying regions in the training set

(A) Case 1: DFT-calculated E_f for designed crystals based on ternary crystals with a target $E_f = -0.5$ eV/atom. The black dashed line is $E_f = -0.5$ eV/atom. Red dots indicate target-satisfying designed crystals. Seven out of 18 (38.9%) inversely designed crystals meet the target E_f within tolerance (gray band of -0.5 ± 0.06 eV/atom).

(B) Case 2: DFT-calculated E_g, and E_f for designed crystals based on ternary and quaternary crystals with target $E_g = 1.5$ eV and $E_f < -1.5$ eV/atom. The gray band indicates the target-satisfying region within tolerance. Red dots indicate target-satisfying designed crystals. Fourteen out of 19 (73.7%) of the inverse-designed crystals have $E_f < -1.5 + 0.06$ eV/atom, and seven (included in the 14) (36.8%) have $E_g = 1.5 \pm 0.3$ eV (0.06 eV/atom and 0.3 eV are user-specified tolerances for E_f and E_g).

(C and D) Histograms of E_g and E_f of the crystals in the dataset of case 2, where target-satisfying crystals only comprise 5.5% of the total dataset (recorded in Table 1), indicating the low likelihood of performing inverse design by pure random chance. The gray bands indicate target-satisfying regions within tolerance. See also Figure S5.

DFT-calculated formation energies are shown in Figure 5A. Seven out of 18 (38.9%) have E_f meeting the target −0.5 eV/atom (within a user-specified tolerance of ±0.06 eV/atom). There are in total eight different crystal structures with more than 30 different elements in those 18 crystals (detailed in section S4.1 in supplemental information), demonstrating that the FTCP framework can access a wide range of structures and chemistries.

To further test the robustness of the framework, we further inversely design another 77 crystals based on ternary crystals with staggered E_f values (−0.3, −0.6, and −0.7 eV/atom). Fourteen out of those 77 solid materials have E_f within the specified tolerance of ±0.06 eV/atom (section S4.2 in supplemental information). Combined with the 18 designed crystals for $E_f = -0.5$ eV/atom, we boost and evaluate the "uniqueness" of these inverse-designed crystal structures in two ways: (1) Composition: we exclude designed crystals whose compositions exist in the Materials Project database. Consequently, all the designed 84 (valid, passed DFT structural relaxation) crystals in this demonstration have unique chemical formulae that do not exist in the Materials Project database. (2) Structure: we investigate if the inverse-designed crystals can be discovered by manual elemental substitution (zero to minimum structural change). We use a structural dissimilarity value, which is calculated based on local coordination information of all sites in the two structures,[33] to assess the structural uniqueness between designed crystals and crystals in the dataset (training + test). The dissimilarity value ranges from zero to above, with zero indicating identical crystal structures, and large values (>1) indicating large dissimilarity. (The dissimilarity value does not consider elements, e.g., comparing NaCl and CsF, which are both $Fm\bar{3}m$, would yield a dissimilarity value of zero.) A low dissimilarity value does not indicate similar crystals per se (such as NaCl and CsF), just similar structures, and thus, we use the dissimilarity value to assess the structural variety of the designed crystals and to suggest the likelihood of discovery by elemental substitution. For our set of 84 valid designed crystals, the median dissimilarity value is 0.37, and 11 crystals have a dissimilarity value above 0.75 (a dissimilarity threshold used in the Materials Project database). The dissimilarity value per designed crystal is taken to be the minimum dissimilarity values between the designed crystal and every crystal in the dataset (training + test), as shown in section S4.3. We posit that designed crystals with high dissimilarity values have a low likelihood of being discovered by conventional elemental substitution.

In the second design case (case 2), we extend the design criteria to multiple objectives by basing the design on ternary and quaternary crystals for bandgap E_g. The design criteria are: (1) bandgap = 1.5 eV (of a user-specified tolerance of 0.3 eV) and (2) formation energy < -1.5 eV/atom (of a user-specified tolerance of 0.06 eV/atom). The bandgap target is selected because it is of interest both for solar cells (near the maximum of the detailed balance efficiency limit for single-junction devices)[31] and for optoelectronic devices (e.g., LEDs). The negative formation energy criterion is chosen as a crude proxy for stability of the designed crystal (see the section synthesizability challenge for a nuanced discussion). We select ternary and quaternary crystals of ≤ 40 sites having energy above hull <0.08 eV/atom in Materials Project as our dataset (training + test). A total of 19 crystals with unique chemical formulae are inversely designed after filtering out compositions that already exist in the Materials Project database (compositional uniqueness). We quantify the median dissimilarity value of designed crystals to be 0.57, and three out of 16 valid crystals have a dissimilarity value above 0.75 (structural uniqueness). The dissimilarity values are shown fully in section S4.3. We perform DFT validation to examine whether these designed crystals' properties meet the user-specified target (after structural relaxation). Figure 5B shows the distribution of E_g and E_f of the 19 designed materials. Seven out of 19 (36.8%) of the designed crystals satisfy the bandgap target $E_g = 1.5\ (\pm 0.3)$ eV, and 14, including the seven, (73.7%) satisfy the formation energy target $E_f < -1.5\ (+0.06)$ eV/atom. We set the user-specified tolerance of E_g of the designed crystals to be significantly larger compared to the one of E_f. This is due to the relatively larger prediction error for E_g compared with E_f in the property mapping of the target-learning branch (Table S3), on which the user-specified tolerances are based. This is attributed to the many zero values of E_g, a situation confusing the property mapping, and leaving the prediction for zero and near-zero values inaccurate. Figures 5C and 5D show the histogram of E_g and E_f values of all crystals in the dataset. The probability of finding crystals that satisfy both $E_g = 1.5\ (\pm 0.3)$ eV and $E_f < -1.5\ (+0.06)$ eV/atom by random sampling (random success rate) is 5.5%, while FTCP reports a success rate of 36.8% (7/19), i.e., a 560% improvement over random, establishing the nontriviality of our inverse design framework.

Case 3: Design for TE power factor (with formation energy and bandgap constraints): Case study for TE applications

In the third design case (case 3), we use the FTCP framework to address one of the outstanding challenges in the field of TEs, i.e., to design new earth-abundant materials that can convert heat to electrical energy and vice versa with high efficiency.[34,35] This design case is challenging from a scientific point of view, because it includes excited-state (as opposed to ground-state) properties, and from an ML point of view, because it relies on sparsely labeled training data (only a small portion of the dataset have TE-relevant property calculated). Example TE-relevant labels include carrier effective mass, Seebeck coefficient, and power factor, which are computationally expensive for DFT or require different computation platforms and are not all present in the same database.[32]

Considering the above, we set the design targets (for designing based on ternary and quaternary crystals) to be: (1) power factor as high as possible (heat-to-electricity conversion efficiency), (2) bandgap between 0.3 and 1.5 eV (desirable for low- and medium-temperature range), and (3) negative formation energy (a preliminary proxy for stability). We select ternary and quaternary crystals of $\leq$40 sites having energy above hull <0.08 eV/atom in Materials Project as our dataset (training + test) for ground-state properties, i.e., E_g and E_f, and we use the database from reference[1] as our dataset (training + test) for power factor, where the constant relaxation time approximation under the Boltzmann transport equations (BTE) is used to calculate the TE-relevant labels. The final dataset has 34,784 crystal structures with ground-state properties (E_g, E_f) from Materials Project. Only 4,284 crystals have corresponding calculated power factor labels from reference.[1]

To tackle the sparse-label problem of the excited-state property, we train a semi-supervised VAE.[36] The semi-supervised VAE allows us to jointly train the dataset with full entries of ground-state properties and partial entries of calculated power factor. The semi-supervised VAE developed in this design case consists of a normal VAE, a target-learning branch that maps its entire latent space to ground-state property labels, and a subset of the latent space to the TE property, power factor. The property-mapping loss of the semi-supervised VAE includes one more component compared with the previous VAE in case 2, i.e., the regression for calculated power factor of 4,284 crystals.

Two crystals out of a total 28 (7.1%) designed crystals are shown in Figure 6. After BoltzTraP calculation, the two design crystals are found to have state-of-the-art power factors, comparable with the best TE materials (for example, GeTe[37] has a similar *c*-axis power factor value to the two FTCP-generated candidate materials). In Figure 6, the doping level and temperature are treated as user inputs. In this design case, the following domain-knowledge-based inverse design criteria are selected: (1) a power factor as large as possible, (2) a bandgap between 0.3 and 1.5 eV (desirable for low- and medium-temperature range), and (3) negative formation energy (a preliminary proxy for stability). We sample the latent space to generate 28 unique crystals after filtering out compositions that exist in the database (compositional uniqueness). The median dissimilarity value of designed crystals is 0.67, and five out of 12 valid crystals have a dissimilarity value above 0.75 (structural uniqueness) shown in section S4.3. (The two designed crystals have dissimilarity values of 0.80 and 0.53.) To examine the designed crystals, we conduct structural relaxation using DFT to obtain the final atomic coordinates and perform BoltzTraP calculations to obtain power factor values.[32]

Opportunities for improvement

Synthesizability challenge

The ultimate goal of inverse design is to pass the designed crystal, supposedly possessing the target property, to an experimental setting, and synthesize the designed

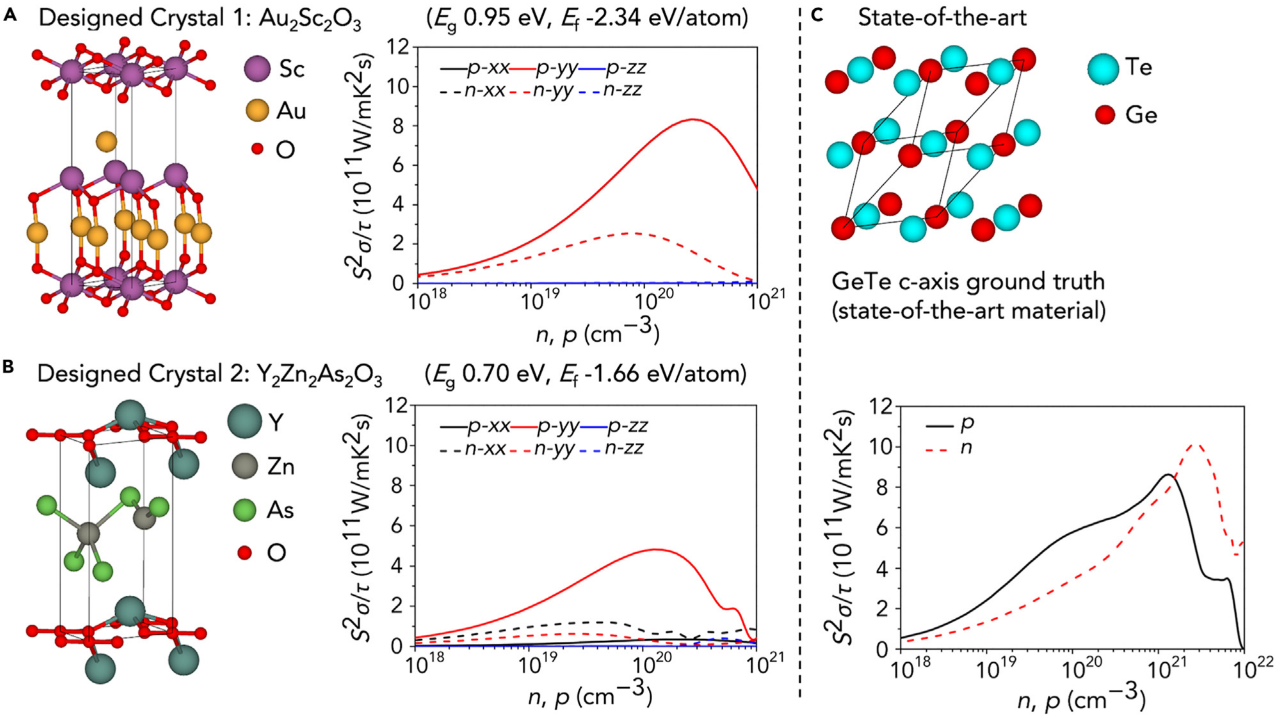

Figure 6. Power factor comparison between two FTCP-designed crystals and the state-of-the-art

(A and B) (A) and (B) show two FTCP-designed crystals with large TE power factors. The two designed crystals are calculated to have comparable power factor (*c*-axis only) with the state-of-the-art high-performance TE material, cubic GeTe, shown in (C). The composition of the two designed crystals do not exist in the Materials Project database (compositionally unique). Their simulated power factor values (divided by the relaxation time) are plotted as a function of doping level at room temperature, both for *n* and *p* doping along the *x*, *y*, and *z* crystal directions in (A) and (B). While E_f and E_g validations were performed using DFT (after structural relaxation), all power factor values were computed under a constant relaxation time approximation using the BoltzTraP package.

crystal. However, the translation from a generative-model-based inverse design algorithm to an actual experimental synthesis of a designed inorganic crystal remains to be demonstrated in the field. More broadly, the topic of synthesis prediction has motivated large research centers (e.g., Center for Next Generation of Materials Design, at https://www.cngmd-efrc.org) indicating both the importance of the topic and the fact that much work remains to be done.

In most prior art to date, instead of experimental validation, proposed inverse design algorithms proffer stability validation using either first-principles calculations,[7,8,14,20–22] or ML surrogate models.[21,23] Among these theoretical stability validations, the majority use negative (or small positive) formation energy as a preliminary proxy for stability (metastability)[8,14,22–24] or small energy above hull as another proxy in the structure-varying cases (limited to specific chemical systems, e.g., V_xO_y).[14,20] The authors also adopt this preliminary proxy of negative formation energy in the design cases (as a design criterion) and ensure the dataset used having a small energy above hull, 0.08 eV/atom. (For FTCP, being both composition and structure varying, using energy above hull poses a challenge because hull diagrams of many chemical systems need to be queried, and especially for quaternary chemical systems, hull diagrams are absent for many chemical systems.) However, a negative formation energy design target, and a small energy above hull in the dataset, are not enough to guarantee synthesizability. In fact, the authors attempted to experimentally validate our model by synthesizing the FTCP-designed $Mn_2Co_2Si_5$ in case 1 (DFT-calculated E_f = -0.326 eV/atom for -0.5 eV/atom target, selected based on synthesis considerations compared with other designed crystals), but upon experimental exploration using an arc melting furnace from pure metal powders, this compound was revealed to decompose into sub-species. A DFT calculation further validates this observation by yielding a decomposition energy of 177 meV/atom.

After this failed experimental validation, the authors realize the importance of optimizing for "precision" (and not "recall" or "F1 score"), framed in language often used in classification tasks. While the FTCP algorithm can design tens of thousands of candidates per hour, experimental throughput is finite. As such, there is a high penalty for false positives. We only need one successful candidate to be experimentally accessible, to succeed at inverse design, but we want to avoid wasting experimental throughput on unsuccessful candidates.

As a potential pathway to future experimental validation, we explore adding a naive synthesizability metric (proxy) named "ICSD score," indicating whether there is an entry for a material in the ICSD. This metric is one method Jang et al. have previously explored[26,27] and often used in drug discovery.[38] The metric is obtained by cross-referencing ICSD and Materials Project databases; those compounds in the Materials Project database with an ICSD entry are labeled "synthesizable (1)," and the rest "not synthesizable (0)." To accommodate the synthesizability metric, we add an output to the target-learning branch to map the latent space to ICSD score. By including the ICSD score loss in the overall property-mapping loss, the latent space organizes to reflect the gradient of both the user-specified target properties and the ICSD score illustrated in Figure 7. During inverse design, although the ICSD score intrinsically contains many false negatives (synthesizable crystals not having ICSD entries), the ICSD score still provides us with information on where to sample in addition to the target properties.[39–43]

The addition of the synthesizability metric also demonstrates the flexibility of the framework to add different metrics during inverse design, should consensus on synthesizability or stability metrics be agreed upon by the field. Besides placing the prediction of the additional metric as an additional output of the target-learning branch, placing a classification model on the designed crystals (after decoding), i.e., filtering the designed crystals instead of the latent points, is also a feasible way. As mentioned previously, much work remains to be completed to develop widely accepted synthesizability metrics. Using the examples presented in this study as guides, the FTCP framework should be able to accommodate future metrics.

Invariance challenge

An aspect of the crystallographic representation that is often considered, arising from the property-prediction applications, is invariance. A crystallographic representation should satisfy full invariances, including translational, rotational, permutational (if there is an order of sites/lattice parameters in the representation), and supercell invariances, because these invariances guarantee different descriptions (such as different CIFs) of the same crystal having the exact same representation, which always results in the same value of the predicted property. A notable class of property-prediction (noninvertible) representation that satisfies full invariances is graph-based, e.g., CGCNN,[28] MEGNet,[3] and SchNet.[44]

However, generative-model-based inverse design (not directed evolution) algorithms need to generate the exact crystal structure (CIF) as the output. It is

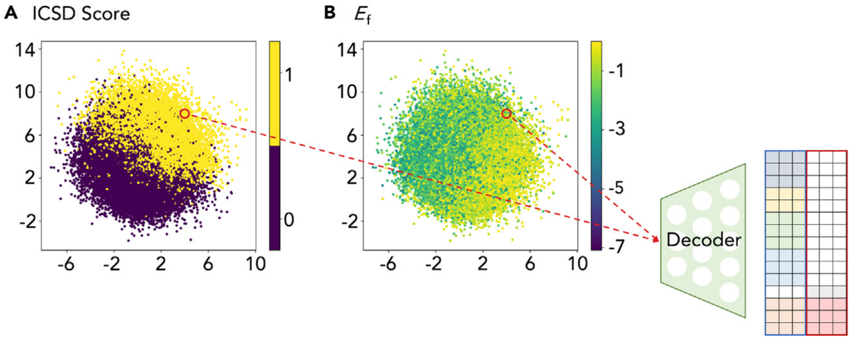

Figure 7. Property-structure latent space with the addition of ICSD score

(A and B) (A) and (B) shows the property-structured latent space colored by the actual ICSD score and the actual formation energy, respectively. Visualization of the property-structured latent space is with the first and the third dimension (256 dimensions in total) of case 1 as an example, after trained with the target-learning branch having an additional output to predict ICSD score. ICSD score is a naive synthesizability metric predicting whether a point in the latent space would have an ICSD entry. Crystals having ICSD entries are labeled 1, indicating "synthesizable", and the rest 0, indicating "not synthesizable." The latent-space colored points shown are crystals in the training set. From the latent space coloring, a separation between the existence and the absence of ICSD entries in (A) is fulfilled on top of the previous E_f property gradient in (B). This benefits the sampling by letting the user sample from the region where the crystals are of higher likelihood to be synthesized (i.e., ICSD scores of 1).

comparably easy to translate the crystal structure to an invariant representation, but the backward translation from the invariant representation to the exact crystal structure is difficult to implement given this is a one-to-many problem. In prior art, we observe that most of the invertible representations do not satisfy any invariances (Table S1) except for the following two (we exclude the composition-only representation[23] from invariance discussion because there is no geometry involved):

1. Crystal site feature embedding (CSFE)[22] satisfies translation invariance. CSFE assumes a 3D grid representation for perovskite structures and fills in the grid with various site/elemental features, which in essence is a composition representation formatted by structure. Since nothing geometrical except for elemental arrangement is included in CSFE, it satisfies translational invariance.
2. Concatenated spectrum representation of composition and powder XRD pattern[24] satisfies full invariances, but it is not fully invertible, i.e., only outputting a powder XRD pattern, which is hard to convert to a crystal unit cell with algorithmic automation.

Therefore, to our knowledge, there are currently no invertible crystallographic representations that satisfy full invariances. Satisfying full/partial invariances for a generalized representation (both composition and structure varying) is an even harder task. This is an interesting and open field of research.

For the current FTCP representation, it also does not satisfy any invariances, determined by the real-space features that solely guarantee invertibility. Kim et al.[20] have shown that performing data augmentation for the real-space representations will lead to a more balanced generation of crystal structures for Mg-Mn-O systems. Given a much larger chemical range (10^5–10^6) is considered in this study, data augmentation is not feasible. The current formulation of reciprocal-space features only preserves permutational invariance to the order of site inputted (atom indexing), but it has the potential to be formulated to preserve more invariances given the invariances in the powder XRD pattern. To provide some quantification, although the FTCP representation is mainly developed to do inverse design, we present the performance degradation used for property prediction/mapping due to translation, rotation, permutation, and different supercells in section S5.

Conclusions and future work

We present a framework for general (both composition- and structure-varying) inverse design of inorganic crystals, called FTCP. This method uses an invertible crystallographic representation comprising concatenated real- and inverse-space features of crystals, where the real-space features are CIF-like, guaranteeing invertibility, and the reciprocal-space features embed periodicity and convoluted elemental properties. By jointly training a VAE with a feedforward target-learning branch, we obtain a probabilistic property-structured latent space that allows for inverse design of crystals with user-specified properties through sampling, decoding, and postprocessing. We use FTCP to perform inverse design and design unique solid-state materials with targeted E_f, E_g, and TE power factor with various chemistries and crystal structures. It is noteworthy that the last property, power factor, is an excited-state property that is challenging to calculate from first principles, yet remains accessible using our inverse design approach through semi-supervised learning. We demonstrate that FTCP can design new crystalline materials that are not in the training set and are dissimilar from known structures. We validate these designed crystals using DFT structural relaxation and confirm their properties by DFT and BoltzTraP calculations. FTCP achieves improvement over random (the probability of finding a material with the user-specified target property by randomly picking from the dataset) ranging from 38.8% to 560%. As a possible pathway to an experimental validation, we explore the addition of a naive synthesizability metric, the existence of an ICSD entry, to further address the synthesizability challenge and demonstrate that FTCP has the flexibility to simultaneously consider this additional metric alongside user-specified target properties. We also comment on the invariance challenge faced by structure-conscious invertible crystallographic representations, including the FTCP representation.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Tonio Buonassisi (buonassi@mit.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

The BTE-calculated power factor is from reference,[1] and the rest of the dataset are queried from the Materials Project[2] in November 2019. (Note a query with the same criteria now would yield a different number of crystals from the recorded number in the study due to the updates and the addition of crystals of the Materials Project.) Source codes, and trained parameters, are available at https://github.com/PV-Lab/FTCP.

Real-space features

To effectively represent a crystalline unit cell in the real space, we extract the necessary information in the CIF, and we concatenate the following information matrices to form CIF-like real-space features, as shown in Figure 2A. (Zero padding is used to satisfy the shape specified when necessary.)

- Element matrix of shape (M, max(N_{elements}, 3)), where M is the length of the one-hot vector representing elements (M = 103 in our case), and N_{elements} is the largest number of components, e.g., $N_{\text{elements}} = 4$ when both ternary and quaternary crystals are in the dataset. The number of columns of the element matrix is set to at least three, to conveniently concatenate to the lattice and site coordinate matrices, which have a minimum of three columns.
- Lattice matrix of shape (2, max(3, N_{elements})). (For the sake of clarity, and without loss of generality, the shape is written as 2 × 3 in section representation.)
- Site coordinate matrix of shape (N_{sites}, max(3, N_{elements})), where N_{sites} is the largest number of sites in the unit cell.
- Site occupancy matrix of shape (N_{sites}, max(3, N_{elements})).
- Elemental property matrix of shape (K, max(N_{elements}, 3)), where K is the length of the elemental property vector $\boldsymbol{Z}$ (K = 92 from atom_init.json of the CGCNN[28] code).

The concatenated real-space features are of shape ($M + 2 + 2{\bullet}N_{\text{sites}} + 1 + K$, max($N_{\text{elements}}$, 3)), where the extra one row of zero padding is added to accommodate the distance of k points in the reciprocal-space features.

Reciprocal-space features

The most common reciprocal-space features of a crystal in materials science are its diffraction pattern. X-ray crystallography is the primary way to study periodic crystals. Modified diffraction images for periodic crystals have been shown to classify their structures accurately.[25] We enrich the information in diffraction images in the reciprocal space by projecting elemental property vector $\boldsymbol{Z}$ to different crystal planes (hkl) with an equation modified from discrete Fourier transformation (Equation 1 and S2.2). The authors find mapping $\mathrm{I}_{hkl} = |\mathrm{F}_{hkl}|^2$ to 2θ as in the powder XRD would impose large sparsity in the data, as only a small number of 2θs actually have signals, and the sparsity problem only becomes worse with a finer grid of 2θ; thus, authors refrain from this mapping and account for it by prepending the distance of the k point (hkl) for each F_{hkl}.

Thus, the reciprocal-space features contain:

- Distance (of k point) matrix of shape (1, $N_{k\text{ points}}$), where $N_{k\text{ points}}$ is the number of nonzero k (hkl) points (59 in our case). $N_{k\text{ points}}$ can be treated as a hyperparameter to be tuned. We obtain $N_{k\text{ points}}$ to be 59 by first limiting $|h| + |k| + |l| \leq 3$ ($N_{k\text{ points}} = 61$) and further tuning ($N_{k\text{ points}}$ reduced to 59).
- FTCP matrix of shape (K, $N_{k\text{ points}}$), where column vector F_{hkl} is arranged according to hkl.

After prepending zero padding of shape ($M + 2 + 2{\bullet}N_{\text{sites}}$, $N_{k\text{ points}}$), we obtain the reciprocal-space features of shape ($M + 2 + 2{\bullet}N_{\text{sites}} + 1 + K$, $N_{k\text{ points}}$).

Generative model

The generative model (VAE) encodes the FTCP representation into a probabilistic latent space of reduced dimension (256) that can be sampled from. An additional target-learning branch mapping the latent vector to material properties further organizes the latent space to reflect continuous variation of the properties/property gradients (thus the name property-structured latent space). Inverse design of new crystals is achieved by sampling different points other than the existing crystals in the property-structured latent space regions that fulfill the user-defined design targets (enabled by the property gradients). Those sampled points are decoded to FTCP-representation-styled outputs using the decoder. With postprocessing, the CIF is then recovered from the real-space features of the outputs.

We treat the reciprocal-space features as a 1D signal with $N_{k\text{ points}}$ channels. The signal in each channel represents the elemental property projection along a specific (*hkl*) direction. We use a 1D convolutional neural network (CNN) to encode the FTCP representation. The 1D CNN in this work is inspired by PointNet used for 3D point sets classification.[20,45] There is a spatial 1D relationship in the reciprocal-space features. Along this spatial axis, our reciprocal-space representations are arranged according to *k* (symmetry) points (*hkl*) that are universal in describing the electronic band structure.

The encoder encodes the FTCP representation into a probabilistic normal distribution ($\mathbf{z}_{\text{mean}}$ with a diagonal covariance matrix, $\mathbf{z}_{\text{variance}}$) using 1D CNN. The decoder, with a symmetrized architecture (using transposed convolutional layers) of the encoder, samples around $\mathbf{z}_{\text{mean}}$ with $\mathbf{z}_{\text{variance}}$ to reconstruct the FTCP representation, and the reconstruction is regularized by the KL divergence between the latent distribution and the standard Gaussian (zero mean and unit variance) prior. In addition to constraining the latent vector distribution to standard Gaussian prior, we also simultaneously train a feedforward target-learning branch to map $\mathbf{z}_{\text{mean}}$ to material properties. The target-learning branch uses Equation 2 where g is fully connected neural networks, σ is the sigmoid function (target properties are normalized to the range of 0–1), and $\mathbf{z}$ is the latent vector. $\mathbf{R}(\mathbf{z})$ is the predicted material properties, and in the case of designing for multiple target properties, instead of a scalar, $\mathbf{R}(\mathbf{z})$ is a vector, of which each entry corresponds to one target material property.

$$R(\mathbf{z}) = \sigma(g(\mathbf{z})) \qquad \text{(Equation 2)}$$

In total, we have three losses:

- $L_{\text{reconstruct}}$, the reconstruction loss, using sum of squares between the reconstructed FTCP representation (matrix) and the inputted FTCP representation (matrix). This is summed over a batch.
- L_{KL}, the KL loss, i.e., the KL divergence. This is averaged over a batch.
- L_{property}, the property-mapping loss, using sum of squares between predicted material properties $\mathbf{R}(\mathbf{z})$ and actual material properties. In case 3, where semi-supervised learning is used, the loss for the incomplete TE label, power factor, is calculated separately also using the sum of squared. (For ICSD score, although a classification task in nature, we still use the sum of squares as loss, taking advantage of the 0–1 range of ICSD score, and the sigmoid activation function.) This is summed over a batch.

The overall loss is

$$L = L_{\text{reconstruct}} + \beta L_{\text{KL}} + \lambda L_{\text{property}} \left(+ \gamma L_{\text{power factor}} \right) \qquad \text{(Equation 3)}$$

where β, λ, and γ are user-defined coefficients. (We give the power factor loss a separate coefficient for the semi-supervised learning.) To learn a disentangled representation, we allow heavy penalization of the latent distribution ($\beta > 1$).[46] We minimize the overall loss with a root mean squared propagation (RMSprop) optimizer.

We train with a batch size of 256, 200 epochs and a dynamic learning rate. The detailed architecture and hyperparameters of the model used in every design case are tabulated in section S2.4 in supplemental information.

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.matt.2021.11.032.

ACKNOWLEDGMENTS

We acknowledge Vladan Stevanovic (Colorado School of Mines) for discussion of applications and validation of FTCP; Shyue Ping Ong (UCSD) for helpful discussions regarding use of ICSD record for stability; Alex Zunger (University of Colorado Boulder) for discussions regarding inverse design; Andy Cooper (Materials Innovation Factory, University of Liverpool) for introducing us to the term "directed evolution" adopted in this manuscript; Aron Walsh (Imperial), Dan Davies (University College London), Keith Butler (Rutherford Appleton Laboratory, UKRI STFC), and Alán Aspuru-Guzik (University of Toronto) for discussions about the upper limits of "stoichiometric inorganic compounds." We appreciate Tian Xie and Xiang Fu (MIT) for the careful review of our manuscript and code. This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) program through the Singapore Massachusetts Institute of Technology (MIT) Alliance for Research and Technology's Low Energy Electronic Systems (LEES) research program. F.O., S.S., and Q.Liang acknowledge support from TotalEnergies SEresearch grant funded through MITei. Y.J. acknowledges the support from the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (No. 2021-0-02068, Artificial Intelligence Innovation Hub). J.L. and X.W. acknowledge support from the Ministry of Education Academic Research Fund R-279-000-532-114,. Y.L. is supported by the National Key Research and Development Program of China (Grant Nos. 2017YFB0702901 and 2017YFB0701502) and the National Natural Science Foundation of China (Grant No. 91641128). S.J. and K.H. acknowledge funding from the Accelerated Materials Development for Manufacturing Program at A*STAR via the AME Programmatic Fund by the Agency for Science, Technology and Research under Grant No. A1898b0043. G.X. is grateful for the support by the Scientific Computing and Data Analysis section of the Research Support Division at Okinawa Institute of Science and Technology Graduate University (OIST). Q.Li is supported by the National Research Foundation (NRF) fellowship grant NRFF13-2021-0106. A.G.A. acknowledges support from Solar EnergyResearch Institute of Singapore (SERIS). SERIS is a research institute at the National University of Singapore (NUS). SERIS is supported by the National University of Singapore (NUS), the National Research Foundation Singapore (NRF), the Energy Market Authority of Singapore (EMA), and the Singapore Economic Development Board (EDB).

AUTHOR CONTRIBUTIONS

Z.R., K.H., and T.B. conceived of this study. Z.R. and S.I.P.T. designed the algorithm and developed and tested the crystallographic representation and the VAE. J.N., G.X., K.H., and Y.J. performed validation tests using DFT and BoltzTraP. J.N. contributed to the analysis of sources of error and role of crystal structure relaxation in correcting for these. F.O. and J.L. contributed to the discussion on the DFT validation and the gap between DFT and experiment. J.L. performed the experiment. S.S., R.Z., and Q.Liang contributed to the development of the synthesizability metric. X.W. and J.L. provided important contributions to improve the neural network (VAE) structure. Q.Li and S.J. provided important intellectual contributions to the section on invariances of crystal representation and shaped the ablation study. Y.L. performed experimental validation and further DFT confirmation on the designed $Mn_2Co_2Si_5$ alloys. A.G.A., Y.L., X.W., Q.Li, S.J., K.H., Y.J., and T.B. supervised different elements of this research. Each co-author contributed portions of the manuscript, the writing of which was coordinated by Z.R., S.I.P.T., S.J., K.H., and T.B. All co-authors approved the final version of the manuscript.

DECLARATION OF INTERESTS

Some of the authors (Z.R., K.H., T.B.) have filed a disclosure on algorithms related to FTCP. Some of the authors (Z.R., K.H., T.B.) are founders and shareholders of a start-up Xinterra, designed to accelerate the development of materials using ML methods.

Received: June 7, 2021
Revised: October 24, 2021
Accepted: November 26, 2021
Published: December 20, 2021; corrected online February 2, 2022

REFERENCES

1. Ricci, F., Chen, W., Aydemir, U., Snyder, G.J., Rignanese, G.-M., Jain, A., and Hautier, G. (2017). An ab initio electronic transport database for inorganic materials. Scientific data *4*, 170085.
2. Jain, A., Ong, S.P., Hautier, G., Chen, W., Richards, W.D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G., and Persson, K.A. (2013). Commentary: the Materials Project: a materials genome approach to accelerating materials innovation. APL Mater. *1*, 011002. https://doi.org/10.1063/1.4812323.
3. Chen, C., Ye, W., Zuo, Y., Zheng, C., and Ong, S.P. (2019). Graph networks as a universal machine learning framework for molecules and crystals. Chem. Mater. *31*, 3564–3572.
4. Zhao, X.-G., Yang, J.-H., Fu, Y., Yang, D., Xu, Q., Yu, L., Wei, S.-H., and Zhang, L. (2017). Design of lead-free inorganic halide perovskites for solar cells via cation-transmutation. J. Am. Chem. Soc. *139*, 2630–2638.
5. Schleder, G.R., Padilha, A.C., Acosta, C.M., Costa, M., and Fazzio, A. (2019). From DFT to machine learning: recent approaches to materials science–a review. J. Phys. Mater. *2*, 032001.
6. Lee, S., and Kang, D.-W. (2017). Highly efficient and stable Sn-rich perovskite solar cells by introducing bromine. ACS Appl. Mater. Interfaces *9*, 22432–22439.
7. Kim, B., Lee, S., and Kim, J. (2020). Inverse design of porous materials using artificial neural networks. Sci. Adv. *6*, eaax9324. https://doi.org/10.1126/sciadv.aax9324.
8. Noh, J., Kim, J., Stein, H.S., Sanchez-Lengeling, B., Gregoire, J.M., Aspuru-Guzik, A., and Jung, Y. (2019). Inverse design of solid-state materials via a continuous representation. Matter *1*, 1370–1384.
9. Noh, J., Gu, G.H., Kim, S., and Jung, Y. (2020). Machine-enabled inverse design of inorganic solid materials: promises and challenges. Chem. Sci. *11*, 4871–4881.
10. Saal, J.E., Kirklin, S., Aykol, M., Meredig, B., and Wolverton, C. (2013). Materials design and discovery with high-throughput density functional theory: the open quantum materials database (OQMD). JOM *65*, 1501–1509.
11. Curtarolo, S., Setyawan, W., Hart, G.L., Jahnatek, M., Chepulskii, R.V., Taylor, R.H., Wang, S., Xue, J., Yang, K., and Levy, O. (2012). AFLOW: an automatic framework for high-throughput materials discovery. Comput. Mater. Sci. *58*, 218–226.
12. Kirklin, S., Saal, J.E., Meredig, B., Thompson, A., Doak, J.W., Aykol, M., Rühl, S., and Wolverton, C. (2015). The Open Quantum Materials Database (OQMD): assessing the accuracy of DFT formation energies. NPJ Comput. Mater. *1*, 1–15.
13. Davies, D.W., Butler, K.T., Jackson, A.J., Morris, A., Frost, J.M., Skelton, J.M., and Walsh, A. (2016). Computational screening of all stoichiometric inorganic materials. Chem *1*, 617–627. https://doi.org/10.1016/j.chempr.2016.09.010.
14. Long, T., Fortunato, N.M., Opahle, I., Zhang, Y., Samathrakis, I., Shen, C., Gutfleisch, O., and Zhang, H. (2021). Constrained crystals deep convolutional generative adversarial network for the inverse design of crystal structures. NPJ Comput. Mater. *7*. https://doi.org/10.1038/s41524-021-00526-4.
15. Franceschetti, A., and Zunger, A. (1999). The inverse band-structure problem of finding an atomic configuration with given electronic properties. Nature *402*, 60–63. https://doi.org/10.1038/46995.
16. Zunger, A. (2018). Inverse design in search of materials with target functionalities. Nat. Rev. Chem. *2*, 0121. https://doi.org/10.1038/s41570-018-0121.
17. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative Adversarial Nets. In Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, *2*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, eds (MIT Press), pp. 2672–2680. https://dl.acm.org/doi/10.5555/2969033.2969125.
18. Diederik, and Welling, M. (2014). Auto-encoding variational bayes. arXiv. arxiv:1312.6114.
19. Asma, N., and Jean, C. (2019). CrystalGAN: learning to discover crystallographic structures with generative adversarial networks. arXiv. arxiv:1810.11203.
20. Kim, S., Noh, J., Gu, G.H., Aspuru-Guzik, A., and Jung, Y. (2020). Generative adversarial networks for crystal structure prediction. arXiv. arXiv:2004.01396.
21. Zhao, Y., Al-Fahdi, M., Hu, M., Edirisuriya, Song, Y., Nasiri, A., and Hu, J. (2021). High-throughput discovery of novel cubic crystal materials using deep generative neural networks. arXiv. arxiv:2102.01880.
22. Choubisa, H., Askerka, M., Ryczko, K., Voznyy, O., Mills, K., Tamblyn, I., and Sargent, E.H. (2020). Crystal site feature embedding enables exploration of large chemical spaces. Matter *3*, 433–448. https://doi.org/10.1016/j.matt.2020.04.016.
23. Dan, Y., Zhao, Y., Li, X., Li, S., Hu, M., and Hu, J. (2020). Generative adversarial networks (GAN) based efficient sampling of chemical composition space for inverse design of inorganic materials. NPJ Comput. Mater. *6*. https://doi.org/10.1038/s41524-020-00352-0.

24. Korolev, V., Mitrofanov, A., Eliseev, A., and Tkachenko, V. (2020). Machine-learning-assisted search for functional materials over extended chemical space. Mater. Horizons *7*, 2710–2718. https://doi.org/10.1039/d0mh00881h.

25. Ziletti, A., Kumar, D., Scheffler, M., and Ghiringhelli, L.M. (2018). Insightful classification of crystal structures using deep learning. Nat. Commun. *9*, 2775.

26. Jang, J., Gu, G.H., Noh, J., Kim, J., and Jung, Y. (2020). Structure-based synthesizability prediction of crystals using partially supervised learning. J. Am. Chem. Soc. *142*, 18836–18843. https://doi.org/10.1021/jacs.0c07384.

27. Sun, W., Dacek, S.T., Ong, S.P., Hautier, G., Jain, A., Richards, W.D., Gamst, A.C., Persson, K.A., and Ceder, G. (2016). The thermodynamic scale of inorganic crystalline metastability. Sci. Adv. *2*, e1600225. https://doi.org/10.1126/sciadv.1600225.

28. Xie, T., and Grossman, J.C. (2018). Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. Phys. Rev. Lett. *120*. https://doi.org/10.1103/physrevlett.120.145301.

29. Martín-Palma, R.J., Martínez-Duart, J., and Agulló-Rueda, F. (2006). Nanotechnology for Microelectronics and Optoelectronics (Elsevier).

30. White, T. (2016). Sampling generative networks. arXiv. arxiv:1609.04468.

31. Shockley, W., and Queisser, H.J. (1961). Detailed balance limit of efficiency of p-n junction solar cells. J. Appl. Phys. *32*, 510–519. https://doi.org/10.1063/1.1736034.

32. Madsen, G.K., and Singh, D.J. (2006). BoltzTraP. A code for calculating band-structure dependent quantities. Computer Phys. Commun. *175*, 67–71.

33. Zimmermann, N.E., and Jain, A. (2020). Local structure order parameters and site fingerprints for quantification of coordination environment and crystal structure similarity. RSC Adv. *10*, 6063–6081.

34. Snyder, G.J., and Toberer, E.S. (2008). Complex thermoelectric materials. Nat. Mater. *7*, 105–114. https://doi.org/10.1038/nmat2090.

35. Recatala-Gomez, J., Suwardi, A., Nandhakumar, I., Abutaha, A., and Hippalgaonkar, K. (2020). Toward accelerated thermoelectric materials and process discovery. ACS Appl. Energy Mater. *3*, 2240–2257. https://doi.org/10.1021/acsaem.9b02222.

36. Kingma, D.P., Rezende, D.J., Mohamed, S., and Welling, M. (2014). Semi-supervised Learning with Deep Generative Models. Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume *22* (MIT Press), pp. 3581–3589. https://dl.acm.org/doi/10.5555/2969033.2969226.

37. Xing, G., Sun, J., Li, Y., Fan, X., Zheng, W., and Singh, D.J. (2018). Thermoelectric properties of p-type cubic and rhombohedral GeTe. Journal of Applied Physics *123* (19). https://doi.org/10.1063/1.5025070.

38. Gao, W., and Coley, C.W. (2020). The synthesizability of molecules proposed by generative models. J. Chem. Inf. Model. *60*, 5714–5723. https://doi.org/10.1021/acs.jcim.0c00174.

39. Griffiths, R.-R., and Hernández-Lobato, J.M. (2020). Constrained Bayesian optimization for automatic chemical design using variational autoencoders. Chem. Sci. *11*, 577–586. https://doi.org/10.1039/c9sc04026a.

40. McDermott, M.J., Dwaraknath, S.S., and Persson, K.A. (2021). A graph-based network for predicting chemical reaction pathways in solid-state materials synthesis. Nat. Commun. *12*. https://doi.org/10.1038/s41467-021-23339-x.

41. Ye, W., Chen, C., Wang, Z., Chu, I.-H., and Ong, S.P. (2018). Deep neural networks for accurate predictions of crystal stability. Nat. Commun. *9*. https://doi.org/10.1038/s41467-018-06322-x.

42. Vasylenko, A., Gamon, J., Duff, B.B., Gusev, V.V., Daniels, L.M., Zanella, M., Shin, J.F., Sharp, P.M., Morscher, A., Chen, R., et al. (2021). Element selection for crystalline inorganic solid discovery guided by unsupervised machine learning of experimentally explored chemistry. Nat. Commun. *12*. https://doi.org/10.1038/s41467-021-25343-7.

43. Kim, E., Huang, K., Saunders, A., McCallum, A., Ceder, G., and Olivetti, E. (2017). Materials synthesis insights from scientific Literature via Text extraction and machine learning. Chem. Mater. *29*, 9436–9444. https://doi.org/10.1021/acs.chemmater.7b03500.

44. Schütt, K.T., Sauceda, H.E., Kindermans, P.J., Tkatchenko, A., and Müller, K.R. (2018). SchNet – a deep learning architecture for molecules and materials. J. Chem. Phys. *148*, 241722. https://doi.org/10.1063/1.5019779.

45. Qi, C.R., Su, H., Mo, K., and Guibas, L.J. (2017). Pointnet: Deep Learning on Point Sets for 3d Classification and Segmentation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (Curran Associates), pp. 652–660. https://openaccess.thecvf.com/content_cvpr_2017/papers/Qi_PointNet_Deep_Learning_CVPR_2017_paper.pdf.

46. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. (2017). Beta-VAE: learning basic visual concepts with a constrained variational framework. ICLR *2*, 6.

Update

Matter

Volume 5, Issue 3, 2 March 2022, Page 1056–1057

DOI: https://doi.org/10.1016/j.matt.2022.01.022

CellPress

Matter

Correction

An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties

Zekun Ren,* Siyu Isaac Parker Tian, Juhwan Noh, Felipe Oviedo, Guangzong Xing, Jiali Li, Qiaohao Liang, Ruiming Zhu, Armin G. Aberle, Shijing Sun, Xiaonan Wang, Yi Liu, Qianxiao Li, Senthilnath Jayavelu, Kedar Hippalgaonkar, Yousung Jung, and Tonio Buonassisi*

*Correspondence:
dannyzekunren@gmail.com (Z.R.),
buonassi@mit.edu (T.B.)

https://doi.org/10.1016/j.matt.2022.01.022

(Matter 5, 336–362, January 5, 2022)

In the Supplemental Experimental Procedures, due to a coding error, the authors mistakenly repeated the boxplot (indicating the statistical distribution of the data) from Figure S6B in the plot of Figure S6C. Of note, only the graphics were incorrect; the values and description in the figure caption was accurate for all cases. The Supplemental Experimental Procedures has now been corrected online. The authors apologize for any confusion.

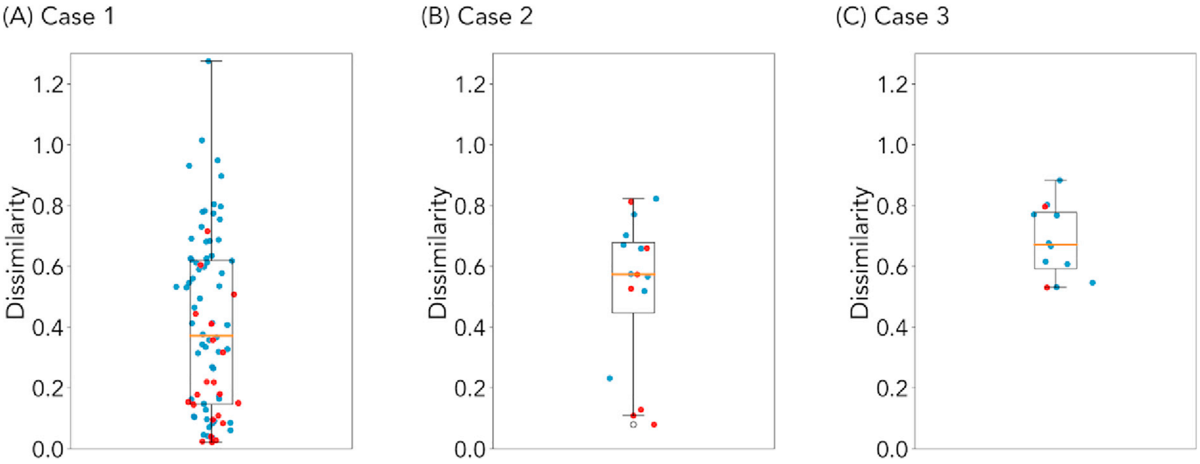

Figure S6. Dissimilarity values of FTCP-designed crystals (corrected)

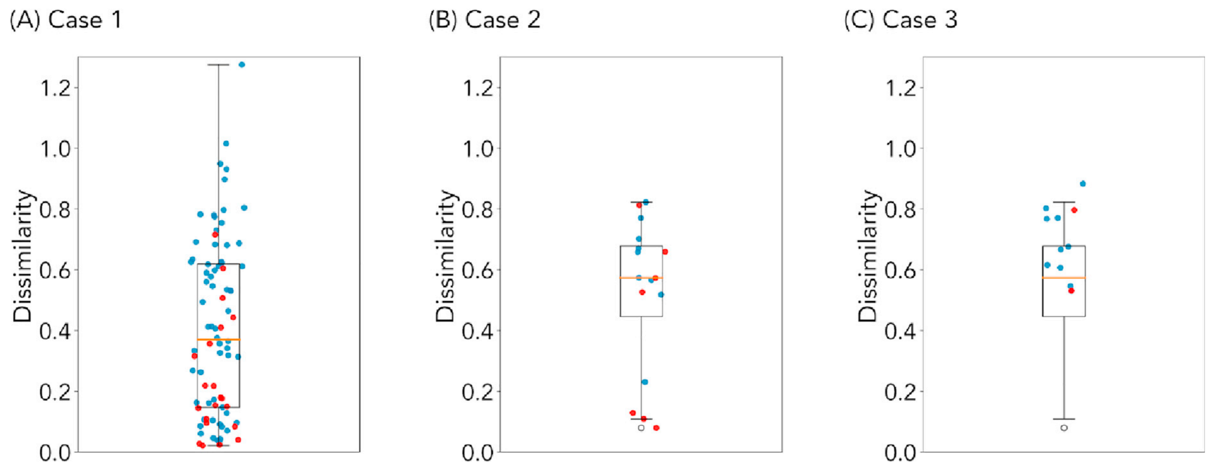

Figure S6. Dissimilarity values of FTCP-designed crystals (original)