Journal Name

ARTICLE TYPE

Cite this: DOI: 00.0000/xxxxxxxxxx

Efficient Symmetry-Aware Materials Generation via Hierarchical Generative Flow Networks

Tri Minh Nguyen [1], Sherif Abdulkader Tawfik [1], Truyen Tran [1], Sunil Gupta [1], Santu Rana[1], Svetha Venkatesh [1]

Received Date
Accepted Date

DOI: 00.0000/xxxxxxxxxx

S1 Geometrical Characteristics

The list of geometrical characteristics of space groups is provided in the Table 1.

S2 Continous GFlownet assumptions

The soundness of the theory of continuous GFlownet relies on a set of assumptions:

- The structure of the state space must allow all states to be reachable from the source state s_0.
- The structure must ensure that the number of steps required to reach any state from s_0 is bounded.
- The learned probability measures need to be expressed through densities over states, rather than over actions.

Our crystal structure state space and training framework satisfy these assumptions as:

- Given the s_0 as the empty crystal structure with space group 1 (lowest symmetry without any constraints on the lattice parameters and no symmetry operation) and initial lattice params, the measurable pointed graph is defined as $\mathscr{S} = s_0 \cup [min_l, max_l]^3 \cup [min_a, max_a]^3 \cup [0,1]^T$ where min_l, max_l are minimum and maximum lattice length, min_a, max_a are minimum and maximum lattice angle.
- We have the upper limit max trajectory length T (defined in Alg. 1) for the trajectory length during the trajectory sampling process. Therefore the number of steps to reach any state is bounded by T.
- The forward p_F and backward p_B policy networks learn the distribution over states $P_F(s_{al}^t | s_{sg}^{t-1}, s_{al}^{t-1}, s_{sg}^t)$, $P_F(s_{sg}^t | s_{sg}^{t-1}, s_{al}^{t-1})$, $P_B(s_{al}^{t-1} | s_{sg}^t, s_{al}^t, s_{sg}^{t-1})$, $P_B(s_{sg}^{t-1} | s_{sg}^t, s_{al}^t)$.

Our crystal structure generation ensures the DAG property of the generation process. Because each step involves adding atoms or refining crystal properties (like the governing space group and lattice parameters), the algorithm cannot cycle back to an identical prior state. The trajectory always moves towards a more complete crystal structure. Even if the same space group number (e.g., 225) is proposed at a later step, the overall crystal state s will be fundamentally different from any previous state where that space group might have been active, as the new state contains a different number of atoms or has modified lattice parameters. The key is that the entire crystal state, $s = (s_{sg}, s_{al})$, evolves, making the construction path acyclic by definition.

S3 Battery material discovery task

Discovering stable material is one of the main targets of machine learning in material science. Finding new battery materials with improved properties such as higher voltage cathodes and higher diffusivity solid-state electrolytes is one of the key challenges in material science [1,2]. We were motivated to introduce a generative approach for battery material discovery which we expect to unravel novel material compositions and structures that are not accessible to high-throughput screening approaches.

Discovering battery material is crucial in many aspects such as safety, environment, and cost. Different from stable material discovery, battery materials are also required to satisfy constraints dedicated to different types of batteries. For example, light-weight, transition-metal-free cation battery materials require a subset of elements. Another example is designing high ionic mobility battery material requires porous structures with low band gap. Such complicated constraints require a flexible model. Our model offers an easy tool to incorporate domain knowledge and requirements into the crystal structure sampling process.

Battery materials, such as cathodes and solid-state electrolytes, can potentially include metallic and heavy elements, even rare-earth elements such as the LLTZ solid-state electrolyte. However, our motivation in this work is to design battery materials composed

[a] Applied Artificial Intelligence Institute, Deakin University. Address: 75 Pigdons Rd, Waurn Ponds VIC, Australia. E-mail: tri.nguyen1@deakin.edu.au, s.abbas@deakin.edu.au, truyen.tran@deakin.edu.au, sunil.gupta@deakin.edu.au, santu.rana@deakin.edu.au, svetha.venkatesh@deakin.edu.au

Table 1 Geometrical characteristics of space groups in terms of lattice angles or lengths

Space group	Crystal family	Lengths constraints	Angles constraints	Parameters search space
1-2	Triclinic c	None	None	$a, b, c, \alpha, \beta, \gamma$
3-15	Monoclinic	None	$\alpha = \beta = 90^\circ$	a, b, c, γ
16-74	Orthorhombic	None	$\alpha = \beta = \gamma = 90^\circ$	a, b, c
75-142	Tetragonal	$a = b$	$\alpha = \beta = \gamma = 90^\circ$	a, c
143-194	Hexagonal	$a = b$	$\alpha = \beta = 90^\circ, \gamma = 120^\circ$	a, c
195-230	Cubic	$a = b = c$	$\alpha = \beta = \gamma = 90^\circ$	a

of Earth-abundant, non-metal elements. There is a large number of research publications that aim at substituting traditional battery materials with metal-free alternatives:

- An anode material that is transition metal-free[3]
- A cathode material that is precious metal-free[4]
- A transition metal-free cathode[5]

In our work, we have limited our attention to the named elements having the highest Earth-abundance according to Table S1 in the seminal work[6]. We have excluded abundant elements such as Mg and Ca, to optimize the computational cost.

S4 Implementation Details

S4.1 State Graph Construction

We determine the edges of the crystal structure graph using k-nearest neighbor atoms within 4 Å. The node feature is the coordinate and the atomic number of the atom.

S4.2 MatErials Graph Network (MEGNet)

MEGNet[7] is a graph neural network that uses two-body distance as edge features. Given the input graph $\mathscr{G} = (\mathscr{V}, \mathscr{E}, u)$ where $\mathscr{X}$ are atoms, $\mathscr{E}$ is edges, and u is global state attribute which is, for example, temperature. In our implementation, the global state attribute is set as 0.0.

First, the bond's attribute (e_k, r_k, s_k) where e_k is the bond feature, r_k and s_k are atoms indices of the bond, are updated as:

$$e'_k = \phi_e[v_{s_k}; v_{r_k}; e_k; u] \tag{1}$$

where ϕ_e is the update function and $[;]$ is concatenation.

Second, the atom feature is updated as:

$$\hat{v_i e} = \frac{1}{N_i^e} \sum_{k=1}^{N_i^e} \{e'_k\}_{r_k=1} \tag{2}$$

$$v'_i = \phi_v[\hat{v_i e}); v + i; u) \tag{3}$$

The update functions ϕ_v, ϕ_e are the MLP layers.

S4.3 Graph Convolution Neural Network (GCN)

GCN[8] is a convolutional network designed to learn the node-level representation of graph structure $\mathscr{G} = (\mathscr{X}, \mathscr{E})$, where $\mathscr{X}$ is the node feature matrix of N nodes and $\mathscr{E} \in R^{N \times N}$ is the adjacency matrix. Given W^l be the weight matrix at l-th layer, the graph convolution operation is then defined as:

$$H^1 = \mathscr{X}, \tag{4}$$

$$H^l = \sigma\left(\tilde{D}^{-\frac{1}{2}} \tilde{\mathscr{E}} \tilde{D}^{-\frac{1}{2}} H^{l-1} W^{l-1}\right), \tag{5}$$

where $\tilde{\mathscr{E}} = \mathscr{E} + \mathscr{I}$ is the adjacency matrix with self-loop in each node. $\mathscr{I}$ is the identity matrix and $\tilde{D} = \sum_j \tilde{\mathscr{E}}_{ij}$ and σ is a non-linear function.

S4.4 Choosing properties to optimize in the reward function

Formation energy, bond length, density, and validity are chosen because they are direct and fundamental properties to be used as the reward function for stable crystal structure generation. Training on these properties allows the model to be universal and flexible enough to be further finetuned to adapt to other tasks or target properties such as band gap.

Some physical interactions such as electronegativity and atomic radius can influence the atoms' position. However, their influence on stability can be indirect and complex. For example, two atoms can still form a bond regardless of the absolute value of the difference in electronegativity (ΔEN) (covalent bond or ionic bond). Atom radius has four common definitions: Van der Waals radius, ionic radius, metallic radius, and covalent radius. Choosing the radius type also depends on the bond types. On the other hand, bond length is a more straightforward and universal measurement.

Terms such as energy above hull E_{hull} which is very informative and directly linked to the stability of the structure. Unfortunately, the computation cost E_{hull} is too high (requires relaxing the structure and then getting the energy of the final state to compute the phase diagram) even using a deep learning model like M3GNet. Therefore, our choice of reward set function also considers balancing between the computation efficiency and fundamentals.

S4.5 Crystal Structure Density

The generated crystal structure density is defined as the ratio of mass m of the unit cell over the volume V of unit cell:

$$P(x) = \frac{m}{V} \tag{6}$$

In our implementation, the density is calcuated using pymatgen library.[9].

S4.6 Diversity Metrics

Structure diversity is computed based on the CrystalNNFingerprint (CNN fingerprint)[10]. CrystalNNFingerprint computes the

fingerprint of a given site i using its coordination features and neighbors. The site i neighbors are determined by the CrystalNN neighbor-finding algorithm. The fingerprint of a crystal structure is the average of fingerprints of all sites.

Composition fingerprint is computed using the statistics of Magpie, computed by element stoichiometry[10,11]. We use ElementProperty.from_preset('magpie') in Matminer as material composition fingerprint.

Then we define diversity as:

$$Diversity = \frac{1}{n} \sum_{i,j \in N_{gen}} d(f_{fp}(i), f_{fp}(j)) \tag{7}$$

where N_{gen} is the set generated crystal structures, n is the number of structure in set N_{gen}. In our experiment, N_{gen} is the top-K crystal structures ranked by reward, f_{fp} is the structure fingerprint in case of structure diversity or composition fingerprint in case of composition diversity, d is the Euclidean distance.

S4.7 Formation Energy Prediction

We use M3GNet[12] to predict the formation energy. As the M3GNet is only trained on the Material Project valid crystal structure, the predictions for invalid structures may be inaccurate and have abnormally low or high formation energy. As our formation energy score function is an exponential function (Sec. 4.2), the formation energy term can get really large, affecting other terms. Therefore, we clamp the formation energy prediction value from -10.0 to 10.0 eV/atom.

S4.8 Hyper-parameters and experiment set up

The hyper-parameter used for our proposed method is presented in Tab. 2. We decide the scaler weights based on the ablation study (Table 3), values over epochs (Fig. 2), and our intuition. In the ablation study, without the bond score term, the result drops significantly. Because the bond score influences the generation quality heavily, we set the weight w_b as 0.5. Both the density and formation energy affect the results at the same level. Therefore, we set both weights w_e and w_p as 0.2. The value of validity over epochs (Figure 2d) converges quickly close to 1.0 which is the maximum value. It seems that validity is an easy term for the model to learn compared to density and formation energy. Therefore, we lower the validity weight w_c to 0.1.

We run our experiments in Tab. 1-4 in the main text three times with random seeds. The result for CDVAE in Table 7 in the main text is reported in their paper.

The training convergence of the SHAFT model, including both the high-level space group policy network (θ_{sg}) and the low-level atom-lattice policy network (θ_{al}), was primarily monitored through several key indicators. A fundamental aspect of training SHAFT is the minimization of the trajectory balance (TB) loss, as defined in Equation 5. Consistent reduction in the TB loss for both policy networks over training epochs was observed in Fig. 1, indicating that the networks were successfully learning to satisfy the flow consistency conditions across the hierarchical state space.

Empirically, the effectiveness of the training and the convergence towards generating desirable structures were also assessed by tracking the average reward of the sampled crystal structures

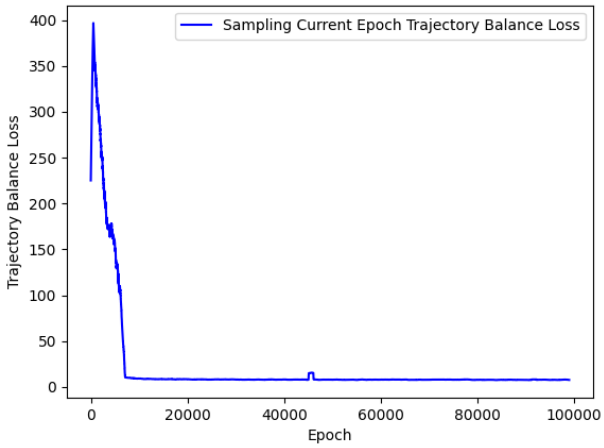

Fig. 1 The trajectory balance loss of SHAFT.

over epochs, as illustrated in Figure 4b in the main text. This figure demonstrates a clear trend of increasing average reward as training progresses, signifying that the model learns to generate structures with increasingly favorable properties according to the defined reward function. Furthermore, Figure 8 shows the evolution of the individual components of the reward function over an extensive number of epochs (10^5 training epochs), providing insight into how the model learns to balance different physical and chemical desiderata. While an exhaustive hyperparameter sensitivity analysis was beyond the scope of this work, the hyperparameters detailed in Table 11 (Appendix A.3) were found to yield stable training dynamics and consistent improvement in both the trajectory balance loss and the quality of generated structures, as evidenced by the aforementioned reward metrics.

Table 2 Hyper-parameters

Hyper-paramters	Value
Learning rate	0.001
Learning rate Z	0.1
Optimizer	Adam
Learning rate scheduler γ	1.0
Initial logZ	0.0
Batch size	32
w_e (Eq. 8)	0.2
w_p (Eq. 8)	0.2
w_b (Eq. 8)	0.5
w_c (Eq. 8)	0.1

S5 MP-Battery dataset

To address the requirements of the battery material discovery task outlined in Sec. 5.1 and Appendix Sec. S3, we developed the MP-Battery dataset using the Materials Project database[13]. We specifically selected materials from the 187,687 available in the dataset that contain light elements such as Be, B, C, N, O, Si, P, S,

and Cl, alongside one of the alkali metals Li, Na, or K. This process resulted in the creation of the MP-Battery dataset, which includes 2,060 distinct materials. The MP-Battery dataset is included in the Supplemental data Github with train/validation/test split.

S6 DFT Validation

We sample and choose the top-K generated materials for DFT verification. Out of 100 crystal structures, 95% (95) are successfully optimized. The complete list of DFT-optimized structures is included in the Supplement data. The distribution of total energy per atom of DFT-optimized structures shows that all DFT- optimized structures have total energy lower than 0 eV/atom.

For the structures that were optimized using DFT, we have computed the crystal decomposition energy (E_D) for each of the structures using the code available in the Python package oganesson[14]. E_D is calculated by obtaining the difference in energy between the structure's energy and the energies of competing phases that are available in Materials Project[13]. For metastable structures (where E_D values are negative), the energy above the convex hull (E_{hull}) is equivalent to $-E_D$. Structures that do not decompose into competing phases will have an $E_H = 0$, and they are known as ground state phases. Out of 100 crystal structures, 7% structure are stable ($E_{hull} = 0eV/atom$) and 56% structure are metastable ($E_{hull} < 0.05eV/atom$).

S7 Reward and Rewards Terms over Epochs

During the training process, we monitor four terms of the reward function: bond score, formation energy, density, and validity score term. The average score terms of each sampled batch over epochs are plotted in Fig. 2.

S8 Distribution of reward terms

The distributions of reward components for 10,000 raw, unfiltered structures generated by each method are plotted in Fig. 3. No top-K filtering has been applied to CD-VAE, SHAFT, GFlownet, Random samples. We also apply top-100 filtering (Filtered Random) using the reward function in Eq. 8 on the random samples to demonstrate the effect of top-K filtering. This provides a direct comparison of the output distributions, demonstrating the quality shift achieved by the trained SHAFT policy compared to the untrained ('Random') model, filtered untrained ('Filtered Random') model, the flat GFlowNet, and CDVAE. Without any training, the sampled structure can be considered random.

S9 Stability evaluated by energy above hull (E_{hull})

The stability of the generated crystal structure can also be evaluated by using the phase diagram to perform convex hull analysis. Convex hull analysis shows the stability of the generated structure composition within the composition space. The energy above hull (E_{hull}) indicates the decomposition energy of the generated composition into the combination of the stable phases. Higher E_{hull} indicates less thermodynamic stability.

Because relaxing structure and calculating energy using DFT is computationally expensive, we can approximate the E_{hull} value using a deep learning model. Specifically, we relax the generated crystal structure using M3GNet Relax tool[12] and then estimate the energy of the final structure using M3GNet[12]. The generated structure's composition and estimated energy are used for convex hull analysis using Pymatgen[9].

S10 Ablation Study

To demonstrate the ability to guide the generative model to generate more stable structures, we perform the ablation study on the reward function's terms of Eq. 8. The generated crystal structures are relaxed using M3GNet optimization framework[12]. In addition, we also compare GCN policy network (details in Appendix Sec.S4.3) and MEGNet policy network (Appendix Sec. 4.3,S4.2). The results reported in Tab. 3 show that all the density, bond score, and formation energy terms are necessary for the model to generate more stable structures. Surprisingly, without the formation energy training signal, the model still can produce structures with negative formation energy with atoms being placed at their preferred distances (distilled from large datasets such as Material Project) combined with appropriate density.

Table 3 The ablation study on the impact of reward function terms on the crystal structure stability. Match rates of the generated crystal paired with structures optimized by M3GNet[12]

Methods	Comp. val.	Struct. val.	Avg. Form. E.	% E. < 0	Match rate ↑	RMS dis. ↓
SHAFT (MEGNet)	1.00	0.99	-1.17	100	0.82	0.16
W/o density	1.00	1.00	-1.20	100	0.76	0.16
W/o bond score	0.99	0.32	-1.43	100	0.53	0.19
W/o formation energy	1.00	1.00	-0.99	96.0	0.77	0.17
W/o validity	1.00	1.00	-1.15	100	0.72	0.17
SHAFT (GCN)	1.00	0.98	-1.15	100	0.81	0.16

The result also shows that it is desirable to have an expressive GNN as a policy network as it can capture the permutation, translation, rotation, and periodic invariance of crystal structure and learn more accurate structure representation.

S11 Complexity Analysis

Our experiments are conducted on a single GPU Tesla V100-SXM2-32GB with 80 cores CPU Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz. We report the average training time per iteration, sampling time per sample, and the number of parameters of policy networks (including backward and forward policy) of our proposed method using the reported configuration and a sampling-based method (CDVAE) in Tab. 4.

We follow the crystal structure graph construction from Mat-

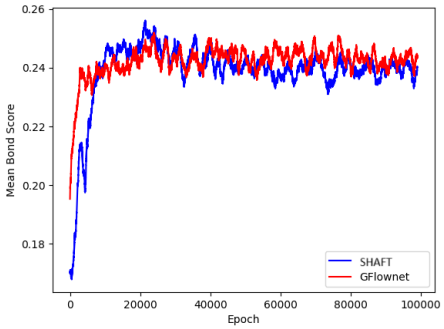

(a) Average bond score over epochs

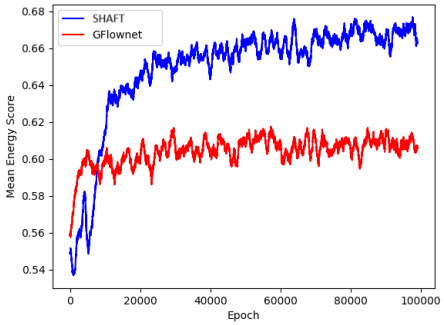

(b) Average formation energy score over epochs

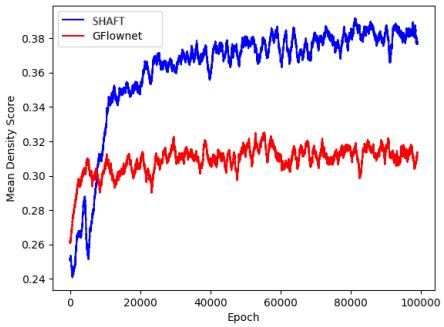

(c) Average density score over epochs

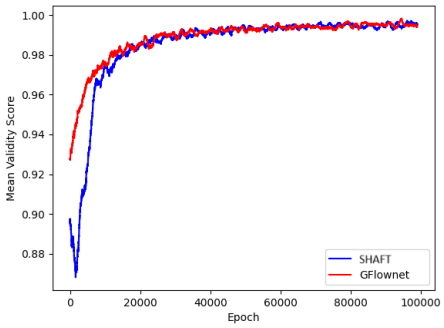

(d) Average validity score over epochs

Fig. 2 Average of terms of reward function over 10^5 epochs.

former[15]. According to complexity analysis[15], assuming each node has at least 12 neighbors, the constructed graph $G = (V, E)$ has $|V| = n$ and $|E| = 6n$. In case of high multiplicity of high symmetry or low-density crystal structure, the number of edges to evaluate can be very high.

We report the number of parameters of our policy networks in Tab. 4. The hierarchical SHAFT model (0.9M parameters) is intentionally more expressive than the flat GFlowNet baseline (0.6M parameters). This larger model capacity is a contributing factor to SHAFT's higher performance in final sample quality, in addition to the exploration efficiency gained from its hierarchical architecture.

Table 4 The average training time per iteration, sampling per sample, and the number of parameters of policy networks of our proposed method using the reported configuration and a sampling-based method (CDVAE)

Methods	Time/epoch	Sampling	Model Params.
CDVAE	3.28s	1.97s	4.9M
SHAFT	0.89s	0.061s	0.9M
GFlownet	0.71s	0.047s	0.6M

References

1 S. Muy, J. Voss, R. Schlem, R. Koerver, S. J. Sedlmaier, F. Maglia, P. Lamp, W. G. Zeier and Y. Shao-Horn, *iScience*, 2019, **16**, 270–282.

2 L. Kahle, A. Marcolongo and N. Marzari, *Energy & Environmental Science*, 2020, **13**, 928–948.

3 J. Zhang, T. Li, Q. Yuan, Y. Wu, Y. Dou and J. Han, *ACS Applied Materials & Interfaces*, 2022, **14**, 54812–54821.

4 A. Riaz, K.-N. Jung, W. Chang, K.-H. Shin and J.-W. Lee, *ACS Applied Materials & Interfaces*, 2014, **6**, 17815–17822.

5 H. Kim, Y.-U. Park, K.-Y. Park, H.-D. Lim, J. Hong and K. Kang, *Nano Energy*, 2014, **4**, 97–104.

6 A. D. Sendek, Q. Yang, E. D. Cubuk, K.-A. N. Duerloo, Y. Cui and E. J. Reed, *Energy & Environmental Science*, 2017, **10**, 306–320.

7 C. Chen, W. Ye, Y. Zuo, C. Zheng and S. P. Ong, *Chemistry of Materials*, 2019, **31**, 3564–3572.

8 T. N. Kipf and M. Welling, Proceedings of the International Conference on Learning Representations, Toulon, France, 2017.

9 S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson and G. Ceder, *Computational Materials Science*, 2013, **68**, 314–319.

10 L. Ward, A. Dunn, A. Faghaninia, N. E. Zimmermann, S. Bajaj, Q. Wang, J. Montoya, J. Chen, K. Bystrom, M. Dylla, K. Chard, M. Asta, K. A. Persson, G. J. Snyder, I. Foster and A. Jain, *Computational Materials Science*, 2018, **152**, 60–69.

11 L. Ward, A. Agrawal, A. Choudhary and C. Wolverton, *npj Computational Materials*, 2016, **2**, 16028.

12 C. Chen and S. P. Ong, *Nature Computational Science 2022 2:11*, 2022, **2**, 718–728.

13 A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards,

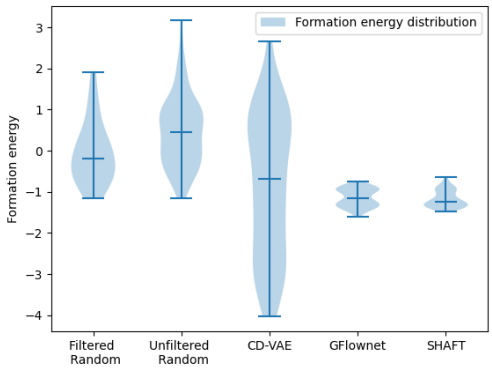

(a) Distribution of formation energy term

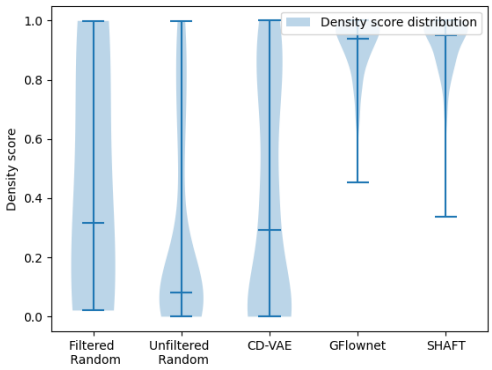

(b) Distribution of density term

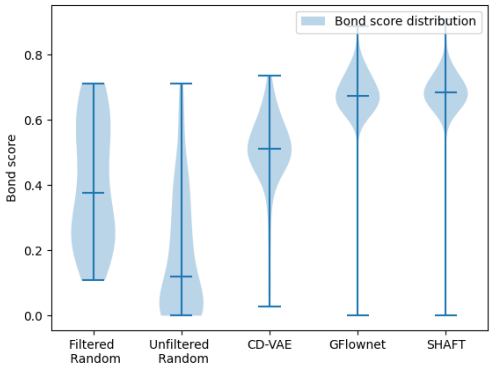

(c) Distribution of bond score term

Fig. 3 The distribution of structures sampled untrained SHAFT (filtered random and unfiltered random), by CDVAE, GFlownet, and SHAFT regarding terms in the reward function.

S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder and K. A. Persson, *APL Materials*, 2013, **1**, 1–12.

14 Sherif Abdulkader Tawfik, *Oganesson https://github.com/sheriftawfikabbas/oganesson*, 2023.

15 K. Yan, Y. Liu, Y. Lin and S. Ji, Proceedings of Advances in Neural Information Processing Systems, 2022.