Apply DFT to DFT: Symmetry-preserved Generation of Crystalline Electron Densities

Yulin Cai [1]

[1]*Cavendish Laboratory, University of Cambridge, United Kingdom, CB3 0HE*. Correspondence to: Yulin Cai yc580@cam.ac.uk.

1. Introduction

Designing materials with desired properties is a task being accelerated by various generative machine-learning approaches, such as variational autoencoder [1], diffusion model [2] and graph neural networks [3]. Many of these techniques perform well for the simulation of interatomic force to guarantee material stability. However, due to the insufficiency of quantum mechanics prior, it is hard to believe such atom-based methods are able to capture subtle electronic structures within the lattice. This makes it challenging to bridge the gap between crystal structure and solid-state properties by noting that a majority of material properties actually result from electron-electron interaction.

According to the Hohenberg-Kohn Theorem, all ground-state properties of a many-electron system are fully determined by its ground-state electron density distribution. In this work, by regarding electron density as the structural representation of solid state materials, we analysed the space-group symmetry of crystal lattices with Fourier transform, and designed a workflow for electron density reconstruction and generation of solid-state materials. The proposed generative model works in the reciprocal domain, where the real-space space-group symmetry is removed to avoid redundancy. Benefiting from careful optimisation in both real and reciprocal space, our method achieves a high level of reconstruction quality in terms of peak signal-noise ratio (PSNR).

2. Apply DFT to DFT

Electron densities vary in shapes and spacial resolution due to their different space groups and probe settings in density functional theory (DFT) calculation. For example, according to the pre-computed data from Materials Project [4, 5], even with the same space group 223, Nb_3Sn is of size (80,80,80) while V_3Si is (72,72,72). These different shapes and sizes need to be aligned carefully, because: 1) the electron density map includes self-repetitive parts due to space-group symmetry, an intriguing property that one cannot expect the machine to learn inherently; 2) inputs for a neural network model generally need to be unified to the same size for forward propagation, noting that simply applying zero-padding at the boundary of smaller cells is not feasible in this scenario as it breaks periodicity.

Applying discrete Fourier transform (DFT) to the electron density map ρ of a primitive cell produced by density functional theory gives

$$\rho(\boldsymbol{G}) = \sum_m \rho(\boldsymbol{r}_m) \exp\left(-\mathrm{i}\boldsymbol{G}\cdot\boldsymbol{r}_m\right)$$

where $\boldsymbol{r}$, $\boldsymbol{G}$ are real and reciprocal grid vectors respectively.

Denote $\{\hat{T}_j\}$ as the set of all symmetry operations of the space group, where $\rho(\boldsymbol{r})$ satisfies

$$\rho(\boldsymbol{r}) = \rho(\hat{T}_j\boldsymbol{r}),\ j = 1, 2, \cdots, N_{\text{op}}$$

Therefore, the reciprocal electron density can be solely represented by the real-space electron densities in the asymmetric unit (AU):

$$\rho(\boldsymbol{G}) = \sum_{m\in\text{AU}} \sum_j \frac{\rho(\boldsymbol{r}_m)}{R_m} \exp\left[-\mathrm{i}\boldsymbol{G}\cdot(\hat{T}_j\boldsymbol{r}_m)\right]$$

where R_m is the repetition counts, i.e., the visiting number of all symmetry operations at point m.

By noting the two summations are commutable, we have

$$\rho(\boldsymbol{G}) = \sum_j \rho_j(\boldsymbol{G})$$

where

$$\rho_j(\boldsymbol{G}) = \sum_{m\in\text{AU}} \frac{\rho(\boldsymbol{r}_m)}{R_m} \exp\left[-\mathrm{i}\boldsymbol{G}\cdot(\hat{T}_j\boldsymbol{r}_m)\right]$$

Since all space-group symmetry operations can be generalised as affine transformations (a rotation and/or a translation), this means in the frequency domain $\rho_j(\boldsymbol{G})$ are phase-transitioned and/or mirror-symmetric to each other. In other words, only with $\rho_1(\boldsymbol{G})$ (where $\hat{T}_1$ is the identical transformation) can we have the full information of the electron density in a primitive cell of a specific space group.

One step further, note that determining $\rho_1(\boldsymbol{G})$ is equivalent to applying DFT to an "almost empty" cell with only the AU being non-zero:

$$\rho_1(\boldsymbol{G}) = \sum_m \rho'(\boldsymbol{r}_m) \exp\left(-\mathrm{i}\boldsymbol{G}\cdot\boldsymbol{r}_m\right)$$

where

$$\rho'(\boldsymbol{r}_m) = \begin{cases} \rho(\boldsymbol{r}_m)/R_m, & m \in \text{AU} \\ 0, & m \notin \text{AU} \end{cases}$$

For electron densities of the same space group with different sizes, the DFTs of their reduced real-space electron densities $\rho'(\boldsymbol{r})$ are of the same size if the rectangular envelope of the maximum AU in the dataset is chosen as the fixed sampling size for the AU of all density maps. Thus, smaller AUs can just be regarded as being zero-padded to the envelope under this aligning strategy.

In order to avoid complex numbers, in practice we adopt discrete cosine transform (DCT) instead of

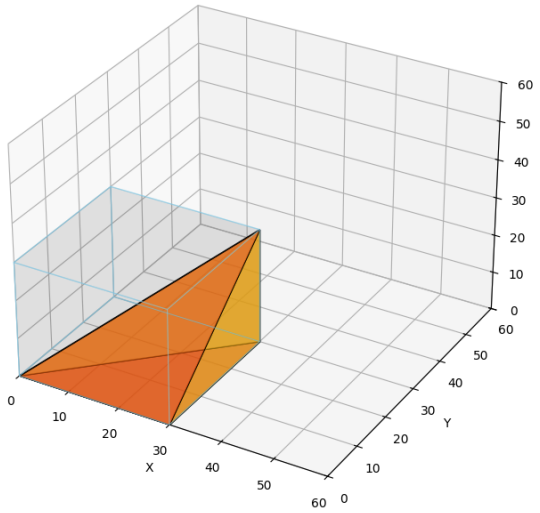

Fig. 1: Visualisation of the asymmetric unit (orange) of the electron density map of $LaNiO_3$ (space group 225). The light rectangular box denotes the area where we apply Fourier transform.

discrete Fourier transform. Since DCT can be regarded as DFT of the even extension of the original signal in 8 octants, we still entitle this section as "Apply DFT to DFT".

3. Workflow

3.1 Data preparation

We collected pre-computed valence electron densities (chgcar) from Materials Project [4, 5] and split them by space groups. As a pioneer attempt, we first conducted experiments on cubic systems, namely, space groups 221, 225 and 227, which have 3,325, 9,448, and 1,437 records respectively. After removing some mistaken samples from the collected data, we divided the data with training:test:validation as 8:1:1 for each space group.

3.2 Loss objective

Since the real-space electron density is a probabilistic distribution in essence, the most natural idea is to minimise the distance between the reconstructed density distribution and the ground-truth distribution, i.e., optimising the Kullback–Leibler (KL) divergence. Simutaneously, we also minimise the mean-square error (MSE) of the reciprocal electron density:

$$\mathcal{L} = \lambda_1 \mathrm{KL}(\rho(\boldsymbol{r}), \rho^*(\boldsymbol{r})) + \lambda_2 \left\|\rho(\boldsymbol{G}) - \rho^*(\boldsymbol{G})\right\|^2 + \lambda_3 \mathrm{KL}_{\mathrm{VAE}}(\boldsymbol{z}, \mathcal{N}(0,1))$$

where $\rho^*(\boldsymbol{r}), \rho^*(\boldsymbol{G})$ denote the ground-truth real and reciprocal electron densities respectively, $\boldsymbol{z}$ is the latent variable, λ_i $(i = 1, 2, 3)$ are tunable weights to average the contribution of each loss type.

3.3 Model

We adopt a variational autoencoder (VAE) [6] with 3D Swin Transformer modules [7] as its encoder and decoder, as shown in Fig. 2.

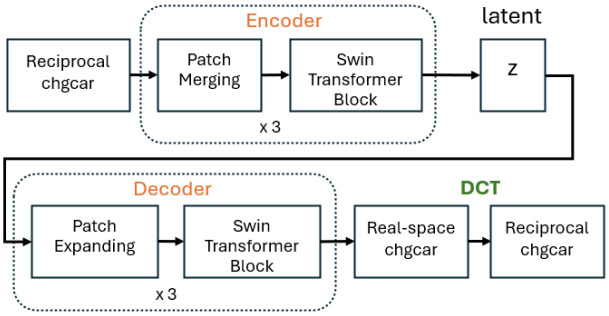

Fig. 2: Pipeline of the Swin-Transformer-based VAE.

The model is trained on 8 NVIDIA A100 GPUs for 300 epochs with batch size 128, using an Adam optimiser with lr=0.001, $\beta = [0.8, 0.999]$. Once trained, the model can be used to generate new chgcars that may represent unseen materials. For generated chgcars that appear in the validation set, we evaluate the average PSNR of the output chgcars in Table 1. In digital image processing, PSNR>40 indicates that the reconstructed image approximately shows no difference in terms of human vision. Although density maps are different from natural images, this MSE-based criterion can still be used as a quantitative indicator. From our results, one can see that the generated electron density distribution shows a high level of quality. Moreover, for generated chgcars that do not appear in the validation set, we are currently applying an improved Bader analysis [8, 9] to determine whether they do represent unseen cubic structures that are convincingly stable.

Table 1: Average PSNR for cubic space groups.

Space Group	Epochs	PSNR
$Pm\bar{3}m$ (221)	300	44.3128
$Fm\bar{3}m$ (225)	300	45.0136
$Fd\bar{3}m$ (227)	300	41.8305

4. Conclusion

In this work, we applied Fourier analysis to capture the asymmetric unit of the electron density distribution of crystal lattices. By fully preserving space-group symmetry in the workflow, we designed an end-to-end generative variational autoencoder to reconstruct and generate electron density maps for the discovery of new solid state materials. Experiments have shown that the reconstruction of crystal electron densities is of high quality, which indicates its potential to make unique discoveries in the unexplored area of targeted space groups. Currently, we are also working on this pipeline to make it applicable for reliable material property prediction.

References

[1] Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. *arXiv preprint arXiv:2110.06197*, 2021.

[2] Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, et al. A generative model for inorganic materials design. *Nature*, pages 1–3, 2025.

[3] Tian Xie and Jeffrey C Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical review letters*, 120(14):145301, 2018.

[4] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.

[5] Kiran Mathew, Arunima K Singh, Joshua J Gabriel, Kamal Choudhary, Susan B Sinnott, Albert V Davydov, Francesca Tavazza, and Richard G Hennig. Mpinterfaces: A materials project based python tool for high-throughput computational screening of interfacial systems. *Computational Materials Science*, 122:183–190, 2016.

[6] Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013.

[7] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.

[8] Graeme Henkelman, Andri Arnaldsson, and Hannes Jónsson. A fast and robust algorithm for bader decomposition of charge density. *Computational Materials Science*, 36(3):354–360, 2006.

[9] W Tang, E Sanville, and G Henkelman. A grid-based bader analysis algorithm without lattice bias. *Journal of Physics: Condensed Matter*, 21(8):084204, 2009.