Global Plane Waves From Local Gaussians: Periodic Charge Densities in a Blink

Jonas Elsborg [1 2] **Felix Ærtebjerg** [1] **Luca Thiede** [3 4]
Alán Aspuru-Guzik [3 4 5] **Tejs Vegge** [1 2] **Arghya Bhowmik** [1 2]

[1]*Department of Energy Conversion and Storage, Technical University of Denmark, DK 2800 Kgs. Lyngby, Denmark* [2]*CAPeX Pioneer Center for Accelerating P2X Materials Discovery, DK 2800 Kgs. Lyngby, Denmark* [3]*Department of Computer Science, University of Toronto, Toronto, Ontario, Canada* [4]*Vector Institute for Artificial Intelligence, 661 University Ave. Suite 710, ON M5G 1M1,15 Toronto, Canada* [5]*Senior Fellow, Canadian Institute for Advanced Research (CIFAR), 661 University Ave.,21 M5G 1M1, Toronto, Canada*. Correspondence to: Jonas Elsborg jels@dtu.dk.

1. Introduction

Kohn–Sham DFT is a central tool for electronic structure calculations, but its self-consistent field (SCF) cycle makes large-scale and high-throughput calculations costly [1, 2]. Machine learning methods can either bypass DFT by predicting observables directly, or preserve DFT fidelity by predicting a high-quality charge density initial guess that reduces SCF iterations [3, 4, 5]. For plane-wave DFT this is particularly natural, yet existing periodic density models often rely on expensive real-space grid evaluations that can negate end-to-end speedups, and a gap exists for fast periodicity-aware density predictors [6, 7, 8, 9].

2. Related work and motivation

Machine-learned charge density prediction enables spatially resolved electronic information and, critically, acceleration of DFT by providing high-quality initial guesses for the SCF cycle. This is increasingly important as DFT remains a major consumer of supercomputing resources and is the backbone of datasets that underpin both task-specific interatomic potentials and emerging foundation models for materials discovery [10, 11, 12, 13]. Existing approaches largely follow two paradigms: orbital models, which expand densities in atom-centered bases and struggle to efficiently capture non-local features in solids [14, 15], and probe models, which predict densities on dense real-space grids using graph neural networks, achieving strong accuracy and SCF reductions but at prohibitive inference cost [8, 9]. While reciprocal space is the natural representation for periodic systems, prior Fourier-domain components for ML charge density modeling have provided limited practical benefit when used only as auxiliary corrections [16]. We therefore introduce the Electronic Tensor Reconstruction Algorithm with Fourier Inversion (ELECTRAFI), which represents periodic charge densities using floating Gaussians with closed-form Fourier transforms, enabling analytic construction of plane-wave coefficients via the Poisson summation formula and a single inverse FFT. ELECTRAFI achieves state-of-the-art accuracy with sub-second inference (up to 633× faster than prior grid-based models) and is the first density model to demonstrate consistent end-to-end DFT wall-time reductions when used for SCF initialization in materials.

3. Methods

We build on the mixture-of-Gaussians density ansatz of Elsborg et al.[17],

$$\tilde{\rho}(\mathbf{r}) = \sum_{j=1}^{N_N} w^{(j)}\, \mathcal{N}(\mathbf{r}; \boldsymbol{\mu}^{(j)}, \Sigma^{(j)}), \tag{1}$$

but reinterpret it for periodic systems. Rather than treating $\tilde{\rho}$ as a physical real-space density, we use it solely as an analytic generator of plane-wave coefficients. Periodicity is enforced via Poisson's summation formula,

$$\sum_{\mathbf{R}\in\Lambda} \rho(\mathbf{r}+\mathbf{R}) = \frac{1}{\Omega} \sum_{\mathbf{G}\in\Lambda^*} \widehat{\rho}(\mathbf{G}) e^{i\mathbf{G}\cdot\mathbf{r}}, \tag{2}$$

which guarantees a smooth periodic density even under reciprocal-space truncation.

The key enabling property is the self-reciprocity of Gaussians under Fourier transformation,

$$\mathcal{F}[\mathcal{N}(\mathbf{r}; \mu, \Sigma)]\,(\mathbf{G}) = \exp\!\big(-\tfrac{1}{2}\mathbf{G}^\top \Sigma \mathbf{G}\big)\, e^{-i\mathbf{G}\cdot\mu}, \tag{3}$$

yielding closed-form plane-wave coefficients

$$\widehat{\rho}(\mathbf{G}) = \sum_{j=1}^{N_N} w^{(j)} \exp\!\big(-\tfrac{1}{2}\mathbf{G}^\top \Sigma^{(j)} \mathbf{G}\big) e^{-i\mathbf{G}\cdot\boldsymbol{\mu}^{(j)}}. \tag{4}$$

A single inverse FFT recovers the periodic real-space density $\rho(\mathbf{r})$, eliminating explicit periodic image summation and dense grid probing required by probe-based models [8, 9]. Electron count is enforced by a global rescaling of $\{w^{(j)}\}$ such that $\int_\Omega \rho(\mathbf{r})\, d\mathbf{r} = N_e$.

Gaussian parameters $(w^{(j)}, \boldsymbol{\mu}^{(j)}, \Sigma^{(j)})$ are predicted locally by a modified transformer-based interatomic potential [18]. Each atom a contributes $n_a = Mv_a$ Gaussians (with valence v_a), whose centers are expressed as displacements from atomic positions, $\boldsymbol{\mu}^{(j)} = \mathbf{R}_{a(j)} + \mathbf{d}^{(j)}$, and whose covariances are parameterized via a Gram factorization $\Sigma^{(j)} = \gamma^{(j)} \mathbf{A}^{(j)} \mathbf{A}^{(j)\top}$. The backbone uses attention-based message passing and dyadic reductions to predict scalar, vector, and tensor channels, avoiding costly high-order SO(3) tensor products while keeping complexity linear in the number of edges. In contrast to direct prediction of $\widehat{\rho}(\mathbf{G})$ [16], ELECTRAFI predicts only local real-space quantities and delegates all global coupling to analytic Fourier structure, aligning the model natively with plane-wave DFT.

Table 1: Charge density prediction on test sets of periodic benchmarks. A dash (—) indicates not reported. Results reproduced from other work are indicated by superscript: [a] [9], [b] [16], [c] [15], [d] [19].

Dataset	Metric	ELECTRAFI	ChargE3Net	DeepDFT	SCDP	GPWNO	InfGCN
MP-Full	NMAE [%]	0.58	**0.54**	0.80[a]	—	—	—
	t_{inf} [s]	**0.24**	78.73	—	—	—	—
	Speedup vs ChargE3Net	**328**×	—	—	—	—	—
GNoME	NMAE [%]	0.93	**0.69**	—	—	—	—
	t_{inf} [s]	**0.15**	33.28	—	—	—	—
	Speedup vs ChargE3Net	**222**×	—	—	—	—	—
ECD-HSE06	NMAE [%]	**1.35**	1.53[d]	—	—	—	—
	t_{inf} [s]	**0.05**	31.65	—	—	—	—
	Speedup vs ChargE3Net	**633**×	—	—	—	—	—
MP-Mixed	NMAE [%]	**1.24**	—	11.50[b]	—	4.83[b]	5.11[b]
	t_{inf} [s]	0.24	—	—	—	—	—
Cubic	NMAE [%]	**1.37**	—	10.37[b]	2.59[c]	7.69[b]	8.98[b]
	t_{inf} [s]	0.08	—	—	—	—	—

Table 2: Results from DFT calculations using various initialization methods for the filtered and recomputed non-magnetic subset of the test files from MP and GNoME. All numbers for time and steps saved are reported relative to the default SAD guess.

Dataset	Metric	SAD (Default)	SC (Converged)	ELECTRAFI	ChargE3Net
MP	NMAE ↓			0.55 %	0.50 %
	ML Init Time ↓			0.24 s	72.11 s
	SCF Steps ↓	16.80	8.73	13.33	12.09
	DFT Time ↓	266.34 s	161.98 s	219.57 s	208.26 s
	Total Time ↓	266.34 s	161.98 s	219.81 s	280.37 s
	DFT Steps Saved ↑		48.04 %	20.65 %	28.03 %
	DFT Time Saved ↑		39.18 %	17.56 %	21.81 %
	Total Time Saved ↑		39.18 %	17.47 %	-5.27 %
GNoME	NMAE ↓			0.88 %	0.59 %
	ML Init Time ↓			0.15 s	28.29 s
	SCF Steps ↓	13.49	6.88	10.11	9.50
	DFT Time ↓	119.98 s	77.91 s	95.06 s	92.67 s
	Total Time ↓	119.98 s	77.91 s	95.21 s	120.96 s
	DFT Steps Saved ↑		51.73 %	25.05 %	29.58 %
	DFT Time Saved ↑		39.75 %	20.77 %	22.76 %
	Total Time Saved ↑		39.75 %	20.64 %	-0.82 %

4. Results

We train and evaluate ELECTRAFI on standard periodic charge density benchmarks using the normalized mean absolute error (NMAE),

$$\text{NMAE} = \frac{\int |\rho_{\text{ref}}(\mathbf{r}) - \rho_{\text{pred}}(\mathbf{r})| \, \mathrm{d}V}{\int \rho_{\text{ref}}(\mathbf{r}) \, \mathrm{d}V}, \tag{5}$$

and report accuracy together with average inference time t_{inf}.

Charge density prediction. Table 1 compares ELECTRAFI against prior density models on periodic benchmarks. ELECTRAFI matches or exceeds state-of-the-art accuracy while reducing inference time by two to three orders of magnitude. On MP-Full and GNoME, ELECTRAFI achieves comparable NMAE to the best prior model, ChargE3Net, with 200–300× faster inference. On the hybrid HSE06 dataset, ECD-HSE06, ELECTRAFI sets state-of-the-art accuracy with a 633× speedup over ChargE3Net. On MP-Mixed and Cubic, ELECTRAFI outperforms all prior models and maintains sub-second inference.

End-to-end DFT acceleration. We assess practical impact by initializing self-consistent DFT calculations with ML-predicted charge densities and measuring reductions in SCF steps and total wall-clock time. Table 2 shows that while ELECTRAFI and ChargE3Net both reduce SCF iterations, only ELECTRAFI yields end-to-end wall-time savings once inference cost is included. Despite slightly smaller SCF step reductions, ELECTRAFI's near-instant inference results in net compute time reductions of ∼ 17–21%, whereas ChargE3Net incurs negative total speedups due to its high inference cost. In Appendix A we show that ELECTRAFI displays favorable scaling for both inference time and DFT acceleration as system size increases. Overall, ELECTRAFI advances the Pareto front of accuracy versus efficiency for periodic charge density modeling by aligning the density representation with the plane-wave formalism of periodic DFT and delegating global structure to analytic Fourier transforms. Currently, end-to-end DFT acceleration benchmarks have yet to mature, but ELECTRAFI's efficiency is immediately useful for downstream settings using charge density as a descriptor, such as molecular dynamics [20] where densities are constructed repeatedly, or in analyses of reactions and chemical bonding based on electronic structure [21].

References

[1] Pierre Hohenberg and Walter Kohn. Inhomogeneous electron gas. *Physical review*, 136(3B):B864, 1964.

[2] Walter Kohn and Lu Jeu Sham. Self-consistent equations including exchange and correlation effects. *Physical review*, 140(4A):A1133, 1965.

[3] Volker L Deringer, Albert P Bartók, Noam Bernstein, David M Wilkins, Michele Ceriotti, and Gábor Csányi. Gaussian process regression for materials and molecules. *Chemical Reviews*, 121(16):10073–10141, 2021.

[4] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1):2453, 2022.

[5] Felix Brockherde, Leslie Vogt, Li Li, Mark E Tuckerman, Kieron Burke, and Klaus-Robert Müller. Bypassing the kohn-sham equations with machine learning. *Nature communications*, 8(1):872, 2017.

[6] Georg Kresse and Jürgen Furthmüller. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Physical review B*, 54(16):11169, 1996.

[7] Paolo Giannozzi, Stefano Baroni, Nicola Bonini, Matteo Calandra, Roberto Car, Carlo Cavazzoni, Davide Ceresoli, Guido L Chiarotti, Matteo Cococcioni, Ismaila Dabo, et al. Quantum espresso: a modular and open-source software project for quantumsimulations of materials. *Journal of physics: Condensed matter*, 21(39):395502, 2009.

[8] Peter Bjørn Jørgensen and Arghya Bhowmik. Equivariant graph neural networks for fast electron density estimation of molecules, liquids, and solids. *npj Computational Materials*, 8(1):183, 2022.

[9] Teddy Koker, Keegan Quigley, Eric Taw, Kevin Tibbetts, and Lin Li. Higher-order equivariant neural networks for charge density prediction in materials. *npj Computational Materials*, 10(1):161, 2024.

[10] Vikram Gavini, Stefano Baroni, Volker Blum, David R Bowler, Alexander Buccheri, James R Chelikowsky, Sambit Das, William Dawson, Pietro Delugas, Mehmet Dogan, et al. Roadmap on electronic structure codes in the exascale era. *Modelling and Simulation in Materials Science and Engineering*, 31(6):063301, 2023.

[11] Stefan Heinen, Max Schwilk, Guido Falk von Rudorff, and O Anatole von Lilienfeld. Machine learning the computational cost of quantum chemistry. *Machine Learning: Science and Technology*, 1(2):025002, 2020.

[12] Maksim Kulichenko, Benjamin Nebgen, Nicholas Lubbers, Justin S Smith, Kipton Barros, Alice EA Allen, Adela Habib, Emily Shinkle, Nikita Fedik, Ying Wai Li, et al. Data generation for machine learning interatomic potentials and beyond. *Chemical Reviews*, 124(24):13681–13714, 2024.

[13] Edward O Pyzer-Knapp, Matteo Manica, Peter Staar, Lucas Morin, Patrick Ruch, Teodoro Laino, John R Smith, and Alessandro Curioni. Foundation models for materials discovery–current state and future directions. *Npj Computational Materials*, 11(1):61, 2025.

[14] Chaoran Cheng and Jian Peng. Equivariant neural operator learning with graphon convolution. *Advances in Neural Information Processing Systems*, 36, 2024.

[15] Xiang Fu, Andrew Rosen, Kyle Bystrom, Rui Wang, Albert Musaelian, Boris Kozinsky, Tess Smidt, and Tommi Jaakkola. A recipe for charge density prediction. *Advances in Neural Information Processing Systems*, 37:9727–9752, 2024.

[16] Seongsu Kim and Sungsoo Ahn. Gaussian plane-wave neural operator for electron density estimation. *arXiv preprint arXiv:2402.04278*, 2024.

[17] Jonas Elsborg, Luca Thiede, Alán Aspuru-Guzik, Tejs Vegge, and Arghya Bhowmik. Electra: A cartesian network for 3d charge density prediction with floating orbitals. *arXiv preprint arXiv:2503.08305*, 2025.

[18] Eric Qu and Aditi Krishnapriyan. The importance of being scalable: Improving the speed and accuracy of neural network interatomic potentials across chemical domains. *Advances in Neural Information Processing Systems*, 37:139030–139053, 2024.

[19] Pin Chen, Zexin Xu, Qing Mo, Hongjin Zhong, Fengyang Xu, and Yutong Lu. Ecd: A machine learning benchmark for predicting enhanced-precision electronic charge density in crystalline inorganic materials. In *The Thirteenth International Conference on Learning Representations*, 2025.

[20] Paolo Vincenzo Freiesleben de Blasio, Peter Bjørn Jorgensen, Juan Maria Garcia Lastra, and Arghya Bhowmik. Nanosecond md of battery cathode materials with electron density description. *Energy Storage Materials*, 63:103023, 2023.

[21] Daniel Koch, Mohamed Chaker, Manabu Ihara, and Sergei Manzhos. Density-based descriptors of redox reactions involving transition metal compounds as a reality-anchored framework: a perspective. *Molecules*, 26(18):5541, 2021.

Appendix A. Scaling and acceleration plots versus system size

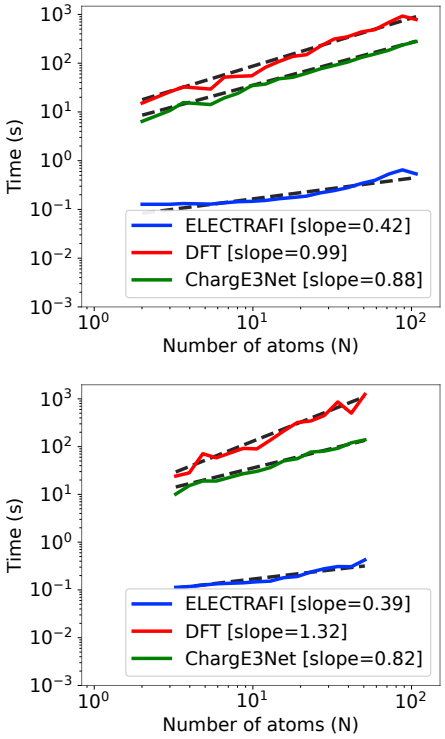

Fig. A1: Logarithmic scaling plots showing wall-clock time versus system size for ELECTRAFI, DFT, and ChargE3Net. **Top:** Materials Project (MP) test set. **Bottom:** GNoME test set.

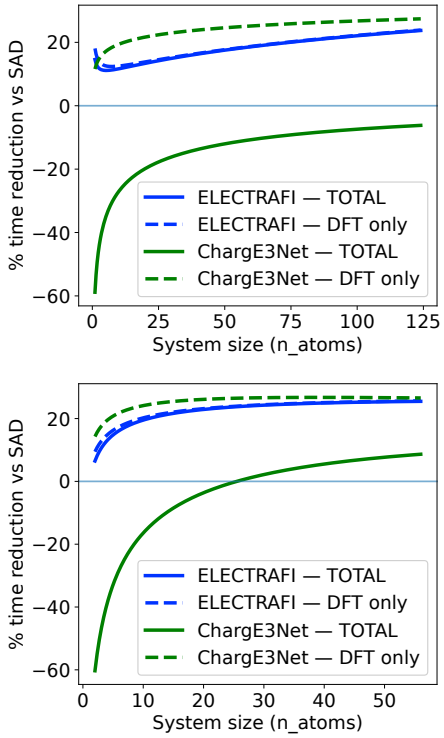

Fig. A2: End-to-end (TOTAL, i.e., including model inference time) vs DFT-only wall-time reduction as a function of system size. **Top:** Materials Project (MP) test set. **Bottom:** GNoME test set.