

# Acceleration of Atomistic NEGF: Algorithms, Parallelization, and Machine Learning

1<sup>st</sup> Mathieu Luisier

*Integrated Systems Laboratory* *Integrated Systems Laboratory* *Integrated Systems Laboratory* *Integrated Systems Laboratory*  
*ETH Zurich* *ETH Zurich* *ETH Zurich* *ETH Zurich*

Zurich, Switzerland

mluisier@iis.ee.ethz.ch

2<sup>nd</sup> Nicolas Vetsch

Zurich, Switzerland

vetschn@iis.ee.ethz.ch

3<sup>rd</sup> Alexander Maeder

Zurich, Switzerland

almaeder@iis.ee.ethz.ch

4<sup>th</sup> Vincent Maillou

Zurich, Switzerland

vmaillou@iis.ee.ethz.ch

5<sup>th</sup> Anders Winka

*Integrated Systems Laboratory* *Integrated Systems Laboratory* *Integrated Systems Laboratory* *Integrated Systems Laboratory*  
*ETH Zurich* *ETH Zurich* *ETH Zurich* *ETH Zurich*

Zurich, Switzerland

awinka@iis.ee.ethz.ch

6<sup>th</sup> Leonard Deuschle

Zurich, Switzerland

dleonard@iis.ee.ethz.ch

7<sup>th</sup> Chen Hao Xia

Zurich, Switzerland

chexia@iis.ee.ethz.ch

8<sup>th</sup> Manasa Kaniselvan

Zurich, Switzerland

mkaniselvan@iis.ee.ethz.ch

9<sup>th</sup> Marko Mladenović

*Integrated Systems Laboratory*  
*ETH Zurich*

Zurich, Switzerland

mmladenovic@iis.ee.ethz.ch

10<sup>th</sup> Jiang Cao

*Integrated Systems Laboratory*  
*ETH Zurich*

Zurich, Switzerland

jiang.cao@iis.ee.ethz.ch

11<sup>th</sup> Alexandros Nikolaos Ziogas

*Integrated Systems Laboratory*  
*ETH Zurich*

Zurich, Switzerland

alziogas@iis.ee.ethz.ch

**Abstract**—The Non-equilibrium Green’s function (NEGF) formalism is a particularly powerful method to simulate the quantum transport properties of nanoscale devices such as transistors, photo-diodes, or memory cells, in the ballistic limit of transport or in the presence of various scattering sources such as electron-phonon, electron-photon, or even electron-electron interactions. The inclusion of all these mechanisms has been first demonstrated in small systems, composed of a few atoms, before being scaled up to larger structures made of thousands of atoms. Also, the accuracy of the models has kept improving, from empirical to fully *ab-initio* ones, e.g., density functional theory (DFT). This paper summarizes key (algorithmic) achievements that have allowed us to bring DFT+NEGF simulations closer to the dimensions and functionality of realistic systems. The possibility of leveraging graph neural networks and machine learning to speed up *ab-initio* device simulations is discussed as well.

**Index Terms**—Device simulation, non-equilibrium Green’s functions, quantum transport, machine learning, GPUs

## I. INTRODUCTION

The continuous shrinking of transistors’ dimensions is posing significant design and fabrication challenges to the semiconductor industry. At the same time, it is bringing the featured sizes of nanoscale devices closer to what *ab-initio* quantum transport (QT) simulators can handle within the framework of the atomistic non-equilibrium Green’s function (NEGF)

formalism, i.e., a few thousands [1]. To reach this number, several progresses have been made since the pioneering works of Taylor et al. [2] and Brandbyge et al., combining density-functional theory (DFT) and NEGF at the beginning of the 2000s. Efficient numerical algorithms have been developed [3], the parallelization of the workload has been widely popularized [4], collaborations with high-performance computing engineers have been established [5], while the emergence of efficient hardware, in particular graphics processing units (GPUs), has paved the way for computer-aided investigations of complex phenomena in nano-transistors [6].

Despite impressive evolution, most *ab-initio* QT simulations are still restricted to the ballistic limit of transport, while electron-phonon or electron-electron interactions might play a detrimental role in nanoscale devices and should therefore be accounted for. Here, we give a brief overview of how the current limitations of QT tools can be addressed, present our implementation of the required models and algorithms into a novel, open source package called *QuaTrEx*, and illustrate it with one example, the simulation of a silicon nano-ribbon in the presence of carrier-carrier scattering. We also suggest how machine learning can be used to bypass the generation of Hamiltonian matrices from first-principles.

## II. AB-INITIO QUANTUM TRANSPORT

Our approach relies on the combination of DFT with NEGF to shed light on the characteristics of nanostructures driven out of equilibrium by an external perturbation (voltage, thermal gradient, or optical signal). The first step consists of creating

This work was supported by the Swiss National Science Foundation (SNSF) under grant n° 209358 (QuaTrEx). We acknowledge support from CSCS under projects lp16 and lp82. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725 (project NEL107).

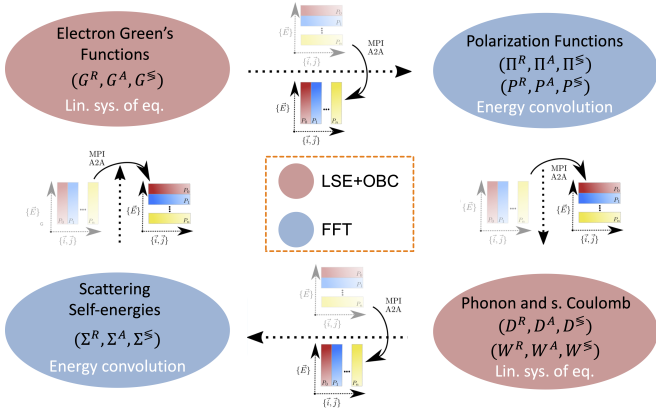


Fig. 1. Iterative scheme to compute the non-equilibrium Green's functions (NEGF) for electrons (retarded:  $G^R$ ) and lesser/greater:  $G^<$  and  $G^>$  in the presence of electron-phonon and/or electron-electron interactions. To determine the corresponding scattering self-energies ( $\Sigma^{R,<}$ ), the phonon ( $D^{R,<}$ ) and screened Coulomb ( $W^{R,<}$ ) Green's functions must be evaluated, which requires the knowledge of the phonon self-energy ( $\Pi^{R,<}$ ) and/or of the polarization function ( $P^{R,<}$ ). All these quantities are matrices  $M$  that depend on the energy ( $E$ ) or frequency ( $\omega$ ). Their entries ( $M_{ij}$ ) represent correlations between two points located at  $R_i$  and  $R_j$ . The calculation of all these matrices involves solving linear systems of equations and open boundary conditions (LSE+OBC, red blocks) or performing energy convolutions through fast Fourier transforms (FFT, blue blocks). In a parallel implementation, it is convenient to store all  $M_{ij}$  entries for one or a few  $E/\omega$  points when dealing with LSE+OBC, while it is preferable to have access to a few  $M_{ij}$  for multiple  $E/\omega$  when computing the  $\Sigma^{R,<}$ ,  $\Pi^{R,<}$ , or  $P^{R,<}$  terms with FFT. Data can be "transposed" with all-to-all communication operations from one representation ( $E, \omega; i, j$ ) to the other ( $i, j; E, \omega$ ).

the Hamiltonian matrix of the system of interest with DFT. Packages relying on a localized basis set, such as CP2K [7] or transformations of plane-wave data into maximally localized Wannier functions (MLWFs) [8], can be utilized for that purpose. The dynamical, bare Coulomb, or electron-phonon coupling matrices can be constructed similarly.

Regardless of the particles (electron, phonon, photon, screened Coulomb) and interactions (electron-phonon, electron-photon, phonon-phonon, electron-electron) considered, the NEGF equations can be written in the following form

$$C^R(\mathcal{E}) = (A(\mathcal{E}) - B^R(\mathcal{E}))^{-1}, \quad (1)$$

$$[C^R(\mathcal{E})]^{-1} \cdot C^>(\mathcal{E}) \cdot [C^R(\mathcal{E})^H]^{-1} = B^>(\mathcal{E}), \quad (2)$$

$$B(\mathcal{E}) \propto \int d\mathcal{E}' (M_1 \cdot C_1(\mathcal{E} - \mathcal{E}') \cdot M_2) \odot C_2(\mathcal{E}'), \quad (3)$$

where  $\mathcal{E} \in \{E, \omega\}$  is the energy ( $E$ ) or frequency ( $\omega$ ) of the Green's function matrix  $C(\mathcal{E})$  and "self-energy"  $B$ , which can be of retarded ( $R$ ), lesser ( $<$ ), or greater ( $>$ ) type. Note that the momentum dependence is omitted for simplicity. The  $C$ ,  $C_1$ , and  $C_2$  matrices are either equal to  $G$  (electron),  $D$  (phonon/photon), or  $W$  (screened Coulomb interaction), while  $B$  refers to the interactions between different physical systems and corresponds to either  $P$  (polarization),  $\Sigma$  (electron scattering self-energy), or  $\Pi$  (phonon/photon scattering self-energy). It also contains an open boundary condition

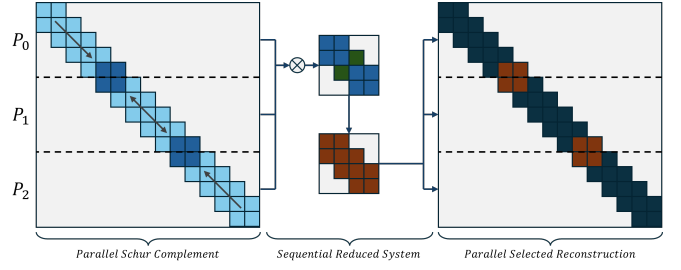


Fig. 2. Simplified representation of the Serinv algorithm [9] for the parallel calculation of the NEGF equations (retarded and lesser/greater). It relies on the repeated application of Schur complements. Starting from a block tri-diagonal input matrix,  $P$  partitions are created, each of them being attributed to different computing units (CPUs or GPUs). After local processing, a reduced system of equations arises, which must be solved sequentially before reconstructing selected entries of the Green's functions locally. Note the mapping between the large and reduced systems (blocks with the same color).

term not shown here. The matrix  $A$  is made of the operator describing the system's intrinsic properties, i.e., of the Hamiltonian, dynamical, or bare Coulomb matrices. Finally, the  $M_1$  and  $M_2$  matrices represent the coupling between different populations (they are equal to the identity matrix in some cases), whereas  $\odot$  indicates an element-wise or standard matrix-matrix multiplication.

Typically, a block tri-diagonal (BT) shape is imposed on all  $A$  and  $B$  matrices by capping the maximum distance over which two particles can interact. As a consequence, only selected entries of  $C$  are needed, those corresponding to the sparsity pattern of  $A$  and  $B$  [1]. From a numerical point of view, Eqs. (1) and (2) are linear systems of equations (LSE) whose outputs are the desired entries of  $C$ . Each  $\mathcal{E}$  is independent of the other and can therefore be treated in parallel. On the other hand, the solution of Eq. (3) involves an energy convolution (EC) where a large number of energies is needed for a few of the  $C$  matrix entries. We have set up a parallel data transposition scheme to adapt the distribution of the multi-dimensional  $B$  and  $C$  tensors in memory, depending on the task at hand (LSE or EC). It is schematized in Fig. 1 in the case of electron-electron and electron-phonon interactions.

Concretely, Eqs. (1) and (2) can be solved with the so-called recursive Green's function (RGF) algorithm, which is sequential in essence [10]. To simulate larger device structures, we have devised a parallel strategy and implemented it into an open-source, GPU-based library, Serinv. It takes advantage of the Schur complement of BT matrices [9] to divide all input/output matrices into multiple partitions. Its principle is summarized in Fig. 2. The energy convolution in Eq. (3) is performed more efficiently in the time domain after a fast Fourier transform (FFT) [11].

### III. RESULTS

As an example, we consider a silicon nano-ribbon with cross-section dimensions similar to experiments [12] (Fig. 3(a)), including electron-electron interactions (self-consistent GW approximation). Its length is first set to 52.1 nm for a total of  $N_A=25,344$  atoms, each of them expressed

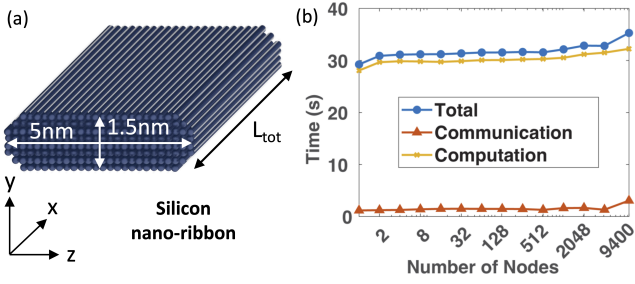


Fig. 3. (a) Schematic view of a Si nano-ribbon with a height of 1.5 nm, a width of 5 nm, and a varying length  $L_{tot}$ . The surface of the nano-ribbon is passivated with hydrogen atoms. (b) Weak scaling results for the Si nano-ribbon in (a) with  $L_{tot}=52.1$  nm and a total of  $N_A=25,344$  atoms, in the presence of electron-electron interactions within the GW approximation, on the Frontier supercomputer. The experiment starts by measuring the execution time on 1 node of the machine with  $N_E=2$  energy points, each of them being handled by 4 GPUs. Subsequently,  $N_E$  increases proportionally with the number of nodes, reaching 18,800 on 9,400 nodes. The total time is decomposed into its communication and computation components.

in a MLWF basis. By leveraging the parallelization approach of Fig. 1 and the distributed NEGF solver of Fig. 2, we could simulate this device up to the full scale (9,400 nodes, 8 GPUs each) of the Frontier supercomputer at Oak Ridge National Laboratory, reaching a weak scaling parallel efficiency of 80% when going from 1 to 9,400 nodes (Fig. 3(b)).

Due to the high computational burden of such calculations, we could not execute a self-consistent Schrödinger-Poisson simulation of a transistor with this nano-ribbon as channel material. Instead, we had to reduce its length to 21.7 nm, keeping the cross section the same, and imposed a 10 nm-long linear potential drop with a 5 nm-long flat region on each side of the structure. The potential difference was set to 0.2 V. Selected results are presented in Fig. 4. From the density-of-states plots in (a), it can be observed that turning on electron-electron interactions (slightly) increases the nano-ribbon band gap, as expected. This limited increase can be explained by the fact that we employ a self-consistent GW scheme, contrary to most DFT packages which implement the  $G_0W_0$  method. Secondly, the electron concentration is high enough (in the order of  $1e16 \text{ cm}^{-3}$ ) to reduce the impact of carrier-carrier scattering on the band gap correction.

The influence of electron-electron interactions on the electronic current distribution is moderate, as can be seen in Fig. 4(b). While the magnitude of the current's peak and its low-energy tail vary along the nano-ribbon transport direction ( $x$ ), they nevertheless remain very close to each other. Still, it is worth noticing that in spite of these variations, the electronic current remains perfectly conserved from one side of the device to the other, validating the implementation of the model, as also demonstrated in [1].

#### IV. OUTLOOK: MACHINE LEARNING

Through parallelization and dedicated algorithms, the solution of the NEGF equations can be massively accelerated, which enables the treatment of large atomic systems at unprecedented levels of accuracy. However, the construction of




Fig. 4. Simulation results for the Si nano-ribbon in Fig. 3(a) with  $L_{tot}=21.7$  nm. A linear potential drop is applied with a voltage difference of 0.2 V between both ribbon extremities. (a) Local density-of-states extracted at  $x=0$  in the ballistic limit of transport (thin blue line) and in the presence of electron-electron interactions within the GW approximation (thick red curve). (b) Spectral current distribution extracted at  $x=2.2$ , 10.9, and 19.5 nm. The inset shows the integral of this quantity (electrical current) as a function of  $x$ , demonstrating current conservation.

the key ingredient for *ab-initio* QT simulations, the DFT Hamiltonian matrix, remains a challenge, especially in structures made of thousands of atoms. Indeed, DFT scales with  $O(N^3)$ ,  $N$  being the number of atoms. Recently, machine-learning techniques based on equivariant graph neural networks (EGNNs) have been proposed to accurately predict the entries of Hamiltonian matrices, after training with DFT data [14]. Such approaches scale with  $O(N)$ , significantly outperforming DFT, but the training cost can be expensive, and most implementations are limited to small domains ( $<150$  atoms).

To be of practical relevance, machine-learned methods should allow for the prediction of multiple Hamiltonian matrices from a single training set. Devices with atomic geometries evolving over time are ideal candidates for that. For example, by applying a voltage to a valence change memory (VCM) cell, oxygen vacancies start moving from one electrode to the other, altering the resistance of this component in the process. To simulate the “current vs. voltage” characteristics of VCMs, their Hamiltonian matrix must therefore be recomputed at each applied voltage. To speed up such calculations, we have developed an EGNN capable of producing the Hamiltonian matrix of devices containing thousands of atoms and have shown that it can return the desired entries for different oxygen vacancy distributions using one single configuration as training [13]. Our methodology is displayed in Fig. 5, whereas the energy resolved transmission function through a VCM cell with a DFT and predicted Hamiltonian is reported in Fig. 6. Although the average error in the machine-learned Hamiltonian entries ( $\sim 2$  meV) is comparable to the state-of-the-art for molecules, it is not yet accurate enough to fully reproduce the behavior of the transmission function.

#### V. CONCLUSION

In this paper, we outlined a parallelization strategy and numerical algorithm tailored to *ab-initio* quantum transport simulations based on DFT and NEGF. Together, they unlock

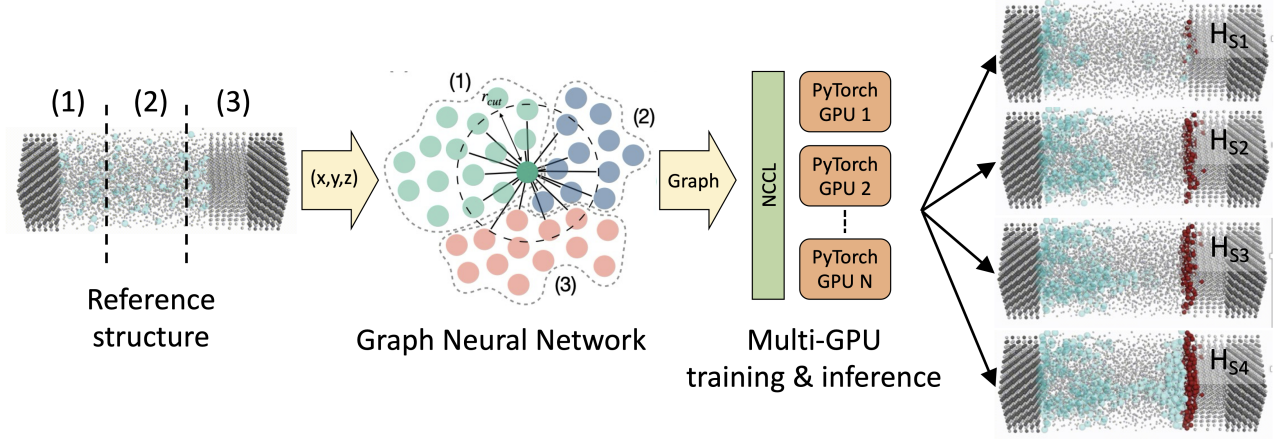


Fig. 5. Illustration on how machine learning can be inserted to accelerate *ab-initio* quantum transport simulations by predicting the entries of Hamiltonian matrices  $H$  based on training from DFT. A single device structure (left), here a valence change memory (VCM) cell, is divided into several partitions and fed to an equivariant graph neural network (EGNN). Once trained with DFT data from the original VCM on multiple GPUs, this EGNN can accurately predict the Hamiltonian matrix of different VCM configurations with various oxygen vacancy distributions. During prediction, the average error of the  $H$  entries does not exceed  $\sim 2$  meV with respect to DFT [13]. Note that EGNNs scale with  $O(N)$ , where  $N$  is the number of atoms, and DFT with  $O(N^3)$ .




Fig. 6. Transmission function through a TiN-HfO<sub>2</sub>-Ti/TiN valence change memory cell as computed from *ab-initio* quantum transport with a DFT Hamiltonian (solid blue curve) and with a Hamiltonian predicted from machine learning after proper training (dashed red curve). The band gap of the HfO<sub>2</sub> insulator layer,  $E_{G,HfO_2}$ , is indicated with a double arrow.

the possibility to explore close-to-experiment device sizes, including relevant physical effects, e.g., carrier-carrier scattering. By further exploiting machine learning, novel opportunities arise, such as the partial elimination of DFT calculations.

#### REFERENCES

- [1] L. Deuschle, J. Cao, A. N. Ziogas, A. Winka, A. Maeder, N. Vetsch, and M. Luisier, "Electron-electron interactions in device simulation via nonequilibrium Green's functions and the GW approximation," *Phys. Rev. B*, vol. 111, p. 195421, 2025. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.111.195421>
- [2] J. Taylor, H. Guo, and J. Wang, "Ab initio modeling of quantum transport properties of molecular electronic devices," *Phys. Rev. B*, vol. 63, p. 245407, Jun 2001. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.63.245407>
- [3] D. E. Petersen, S. Li, K. Stokbro, H. H. B. Sørensen, P. C. Hansen, S. Skelboe, and E. Darve, "A hybrid method for the parallel computation of Green's functions," *Journal of Computational Physics*, vol. 228, no. 14, pp. 5020–5039, 2009. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0021999109001648>
- [4] G. Klimeck, "Parallelization of the nanoelectronic modeling tool (NEMO 1-D) on a beowulf cluster," *Journal of Computational Electronics*, vol. 1, no. 1, pp. 75–79, Jul 2002. [Online]. Available: <https://doi.org/10.1023/A:1020767811814>
- [5] S. Li, S. Ahmed, G. Klimeck, and E. Darve, "Computing entries of the inverse of a sparse matrix using the FIND algorithm," *Journal of Computational Physics*, vol. 227, no. 22, pp. 9408–9427, 2008.
- [6] A. N. Ziogas, T. Ben-Nun, G. I. Fernández, T. Schneider, M. Luisier, and T. Hoefer, "A data-centric approach to extreme-scale ab initio dissipative quantum transport simulations," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '19, 2019. [Online]. Available: <https://doi.org/10.1145/3295500.3357156>
- [7] T. D. Kühne *et al.*, "CP2k: An electronic structure and molecular dynamics software package - quickstep: Efficient and accurate electronic structure calculations," *J. Chem. Phys.*, vol. 152, p. 194103, 2020.
- [8] N. Marzari and D. Vanderbilt, "Maximally localized generalized Wannier functions for composite energy bands," *Phys. Rev. B*, vol. 56, pp. 12 847–12 865, 1997. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.56.12847>
- [9] V. Maillou, L. Gaedke-Merzhaeuser, A. N. Ziogas, O. Schenk, and M. Luisier, "Serinv: A scalable library for the selected inversion of block-tridiagonal with arrowhead matrices," 2025. [Online]. Available: <https://arxiv.org/abs/2503.17528>
- [10] A. Svizhenko, M. P. Anantram, T. R. Govindan, B. Biegel, and R. Venugopal, "Two-dimensional quantum mechanical modeling of nanotransistors," *J. Appl. Phys.*, vol. 91, pp. 2343–2354, 2002. [Online]. Available: <https://doi.org/10.1063/1.1432117>
- [11] K. S. Thygesen and A. Rubio, "Conserving GW scheme for nonequilibrium quantum transport in molecular contacts," *Phys. Rev. B*, vol. 77, p. 115333, 2008. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.77.115333>
- [12] A. Agrawal *et al.*, "Silicon ribbonfet CMOS at 6nm gate length," in *2024 IEEE International Electron Devices Meeting (IEDM)*, 2024, pp. 1–4. [Online]. Available: <https://ieeexplore.ieee.org/document/10873367>
- [13] C. H. Xia, M. Kaniselman, A. N. Ziogas, M. Mladenović, R. Mahjoub, A. Maeder, and M. Luisier, "Learning the electronic Hamiltonian of large atomic structures," 2025. [Online]. Available: <https://arxiv.org/abs/2501.19110>
- [14] H. Li, Z. Wang, N. Zou, M. Ye, R. Xu, X. Gong, W. Duan, and Y. Xu, "Deep-learning density functional theory Hamiltonian for efficient ab initio electronic-structure calculation," *Nature Computational Science*, vol. 2, no. 6, p. 367–377, 2022. [Online]. Available: <http://dx.doi.org/10.1038/s43588-022-00265-6>