

Deep learning extraction of band structure parameters from density of states: A case study on trilayer graphene

Paul Henderson ¹, Areg Ghazaryan,² Alexander A. Zibrov,^{3,4} Andrea F. Young,⁴ and Maksym Serbyn ²

¹*School of Computing Science, University of Glasgow, Scotland*

²*Institute of Science and Technology Austria (ISTA), Am Campus 1, 3400 Klosterneuburg, Austria*

³*Department of Physics, Harvard University, Cambridge, Massachusetts, USA*

⁴*Department of Physics, University of California, Santa Barbara, California, USA*



(Received 17 May 2023; accepted 22 August 2023; published 11 September 2023)

The development of two-dimensional materials has resulted in a diverse range of novel, high-quality compounds with increasing complexity. A key requirement for a comprehensive quantitative theory is the accurate determination of these materials' band structure parameters. However, this task is challenging due to the intricate band structures and the indirect nature of experimental probes. In this work, we introduce a general framework to derive band structure parameters from experimental data using deep neural networks. We applied our method to the penetration field capacitance measurement of trilayer graphene, an effective probe of its density of states. First, we demonstrate that a trained deep network gives accurate predictions for the penetration field capacitance as a function of tight-binding parameters. Next, we use the fast and accurate predictions from the trained network to automatically determine tight-binding parameters directly from experimental data, with extracted parameters being in a good agreement with values in the literature. We conclude by discussing potential applications of our method to other materials and experimental techniques beyond penetration field capacitance.

DOI: [10.1103/PhysRevB.108.125411](https://doi.org/10.1103/PhysRevB.108.125411)

I. INTRODUCTION

Electronic band structure of crystalline solids represents a simple yet very rich example of emergence. Under the influence of scattering from the lattice potential, the electron may acquire a different value of effective mass, become massless, and acquire additional quantum numbers such as pseudospin. In addition, electronic band structure determines basic properties of materials, provided interactions are weak enough [1]. Therefore, identifying material parameters that determine the band structure is of crucial importance. From a theoretical point of view *ab initio* methods such as density functional theory have achieved enormous success in this direction [2]. Nevertheless one typically relies on experimental data to quantitatively extract band structure properties. Experimentally there exist numerous ways to access the electronic structure, such as angle-resolved photoemission [3] and x-ray absorption spectroscopy [4], de Haas-van Alphen effect based on magnetic oscillations [5], analyzing reflection and absorption spectra [6], and electronic transport measurement [7], to name just a few. Despite such a wealth of measurement techniques, matching experimental results with theoretical predictions remains a challenging problem due to the complexity of the band structure, which translates into a large number of involved parameters.

The recent surge of two-dimensional (2D) materials [8] brings new aspects to the problem of determining band structure. First, often the complexity and the number of parameters in 2D materials is considerably lower compared to their three-dimensional counterparts. Besides, 2D materials feature additional level control such as modifying charge density by

gating, and may have an extremely high crystal quality. This opens access to high-resolution experimental data, which may potentially be used for precise determination of band structure parameters. A particular example of such data is provided by so-called penetration field capacitance measurements [9], that effectively probe the density of states (DOS) of the material as a function of carrier density and transverse electric field. Such experimental data has been used to determine material parameters such as hopping matrix elements in several 2D systems [10,11].

Typically, extraction of band structure parameters based on experimental data relies on an efficient solution to what we term the forward problem. In the specific example of penetration field capacitance measurements sensitive to the DOS, this means simulating the DOS for specific values of material parameters such as hopping matrix elements entering tight-binding model of the band structure. However, the existence of an efficient solution for this forward problem does not guarantee a fast solution to the inverse problem—identification of the physical parameters corresponding to a set of empirical data. The inverse problem is challenging because (i) solving for the best-fitting parameters is a high-dimensional optimization problem that requires numerous simulations of the forward problem at each step that can quickly become very costly numerically; and (ii) experimental measurements are typically affected by additional factors not easily accounted for in simulation (e.g., geometric and parasitic capacitance, disorder), meaning that an exact match between the data simulated in the forward problem and that obtained from experiments is not possible. The typical approach is therefore manual comparison of an experimental dataset with a large number of

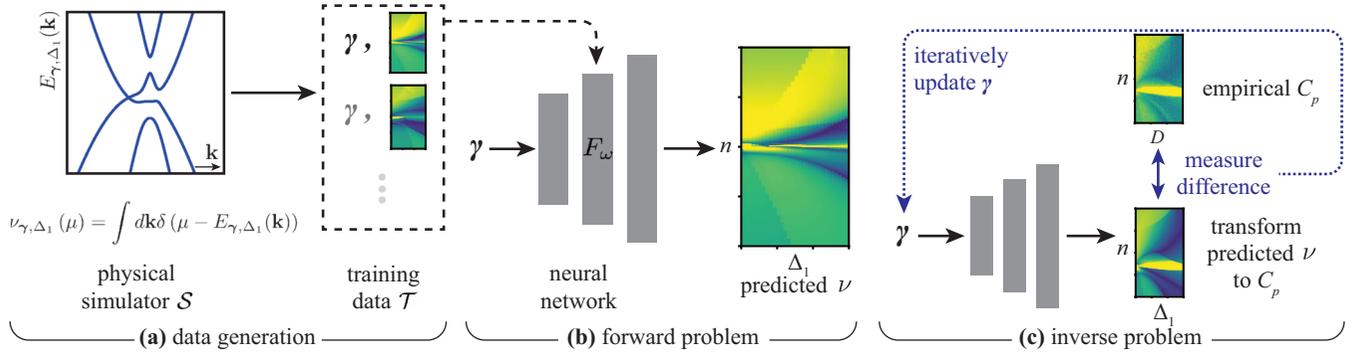


FIG. 1. Our approach to the direct extraction of the band structure parameters uses datasets obtained with a physical simulator shown in (a) to train a DNN in (b), thus providing a more efficient solution to the forward problem. The DNN-based simulator is used for a gradient optimization of the band structure parameters, allowing to extract their values from experimental data in (c).

simulated ones, relying on physical intuition of which features are important. This process is laborious and computationally expensive,¹ calling for the development of more efficient and systematic approaches.

In this work we present a machine-learning-based method that automates the process of comparing numerical simulation and experimental data, so the physical parameters of the band structure that gave rise to a particular experimental dataset can be determined with minimal human effort. Recently machine learning and artificial neural network techniques have seen various applications in the realm of physical sciences [12]. In condensed matter physics, artificial neural networks have been used to represent quantum states [13,14] and learn these states from available data [15,16]. In a different direction, recently machine learning models were used for photonic crystals band diagram prediction and gap optimization [17–19]. Despite a large number of more theoretical applications, machine learning approaches are only starting to be employed in analysis of experimental data. Recent examples include identification of quantum phase transitions [20] and hidden orders from experimental images [21]. These few examples highlight the strong potential of machine-learning-based approaches on experimental data, that we further exploit in the present work.

A conceptual overview of our approach is shown in Fig. 1. To extract the band structure parameters from experimental data, we first train a deep neural network (DNN) [22,23] that solves the forward problem by replicating the numerical calculation of the DOS (Sec. II A). To this end we use the simulation of the experimental data shown in Fig. 1(a). In the particular example of penetration field capacitance data considered here, the simulator uses the band structure parameters, the asymmetry potential between two edges of the system (physically equivalent to transverse electric field) and the chemical potential as input parameters. As an output we get charge density and from that determine the DOS by differentiating density with respect to chemical potential. A set of simulated data is used to train the DNN in Fig. 1(b). Constructed in a way to efficiently replace the data simulator, the DNN acts as a function that takes the band structure parameters, the asymmetry potential and directly charge density as

input, and outputs the corresponding DOS. It is constructed by learning from a large dataset of simulation results, optimising its output to always match that of the simulator. The resulting DNN represents a fast and differentiable replacement for the physical simulation. It can therefore be used to efficiently solve the inverse problem (Sec. II B). In particular, the values of parameters that gave rise to a given dataset are extracted using gradient-based optimization in Fig. 1(c), where we iteratively modify the band structure parameters until the DNN’s output matches the provided DOS values.

The task of mapping a vector of inputs (e.g., band structure parameters) to a continuous output (e.g., DOS) is known as regression in machine learning [23]. A DNN implements such a mapping as a series of chained matrix multiplications (layers) interleaved with elementwise nonlinear functions (activations). Each layer multiplies the vector of outputs from the previous by some weight matrix, to give an updated vector [22]; the final layer typically yields a single value. The weight matrices are optimized (trained) using first-order optimization (e.g., gradient descent), such that the overall mapping from inputs to output approximates some function defined by a training dataset of inputs and desired outputs. The celebrated universal approximation theorem [24] proves that a neural network with just two layers (but unbounded width) can represent any smooth function. More recently, it has been shown that this is also true for a neural network of bounded width (but unbounded depth) [25]. In practice, even finite-sized DNNs have proven very successful in approximating complex functions in many domains of science. One of our contributions is to show a DNN can also provide an accurate estimate of DOS given band structure parameters, field strength, and chemical potential.

Since our final goal is to determine the band structure parameters from experimental measurements of penetration capacitance, it might seem natural to train the DNN for exactly this task (the inverse problem), instead of the forward problem. However, this is not possible in practice. We have only a single experimental dataset, for which the parameters are unknown, whereas machine learning techniques require a large dataset of training examples (with the true output known) to learn the desired mapping from. If instead we trained on easy-to-obtain simulated data, the resulting model would not work on experimental data since the latter is

¹This approach is aptly referred to as gradient descent.

significantly different from the former both in terms of the relative magnitude of features in the data, and the locations of features such as edges. These differences may arise since the simulator uses a simplified effective model of the material and does not account for screening at the microscopic level, disorder, strain, experimental uncertainties, and possibly other ingredients. In contrast, statistical learning theory requires that the training and test data be drawn from the same distributions if a trained model is to work on the latter [26].

As a specific application, we demonstrate the framework outlined above on Bernal stacked (ABA) trilayer graphene. For this material both band structure parametrization [27,28] and high-quality experimental measurements are readily available [11], calling for an accurate extraction of the band structure parameters. The determination of the band structure was performed by *tour de force* manual fitting in an earlier work [11], thus allowing us to benchmark our approach. First, we train the DNN and show it gives an efficient and accurate surrogate for numerical calculation of the DOS for this system, for a wide range of band structure parameters (Sec. III C). Next, we use the DNN for automatically solving the inverse problem of determining the physical parameters giving rise to certain values of the penetration capacitance (Sec. III D). Finally, we apply this to experimental data from Ref. [11] by exploiting techniques from computer vision, that allow matching important features of the measurements (e.g., Van Hove singularity peaks and jumps of DOS), while ignoring features that differ between experimental and simulated data (e.g., measurement noise and discreteness of calculation grid) (Sec. III E). The resulting values of parameters agree within error bars with the estimates from the literature, thus providing a particular benchmark for our approach.

The paper is organized as follows. Section II describes the general structure of the DNN for the forward problem and our minimization approach for inverse problem. Section III applies the framework constructed in Sec. II to ABA graphene simulation results for the forward and inverse problems, eventually utilizing it for extracting band parameters from experimental results. Finally Sec. IV is devoted to discussion of the main results and generalization of the model to other systems.

II. METHOD

We assume access to a simulator \mathcal{S} (see Fig. 1) that in the particular example of the penetration field capacitance calculates the charge density n and density of states $\nu = \frac{\partial n}{\partial \mu}$ given band structure parameters $\boldsymbol{\gamma}$, interlayer asymmetry Δ_1 , and chemical potential μ . We assume $\boldsymbol{\gamma} \in \Gamma$, where Γ defines a physically-plausible range for those parameters, and similarly that $(\Delta_1, n) \in P$. The approach presented here is general, while the specific physical meaning of these parameters will be discussed in Sec. III. Typically the simulator will be slow to evaluate, making it difficult to use in the loop for solving the inverse problem, of finding the physical parameters $\boldsymbol{\gamma}$ corresponding to observed data. We will instead use \mathcal{S} to generate training data represented by tuples $(\boldsymbol{\gamma}, \Delta_1, n, \nu)$ for a machine learning regression model—a deep neural network— F_ω , that will be trained to approximate \mathcal{S} in Sec. II A. We will then

use the resulting DNN F_ω when solving the inverse problem in Sec. II B.

A. Forward problem

We introduce a function F_ω that maps $\boldsymbol{\gamma}$, Δ_1 and n to the DOS. We choose F_ω to be a deep neural network (DNN) [22], with weights ω ; these weights are free parameters that determine the function it represents. Our goal is that F_ω matches \mathcal{S} as closely as possible for all relevant values of $\boldsymbol{\gamma}$, Δ_1 and μ , i.e., if the simulator returns n and ν for given $(\boldsymbol{\gamma}, \Delta_1, \mu)$ and if $(\Delta_1, n) \in P$ are within the domain of physically realistic parameters, then DNN approximates well the DOS, $F_\omega(\boldsymbol{\gamma}, \Delta_1, n) \approx \nu$. The network weights ω would ideally be set to minimize the absolute difference between the network's predicted values and those ν from the simulator, over the entire parameter space $\Gamma \times P$:

$$\omega = \operatorname{argmin} \omega' \int_{(\Delta_1, n) \in P, \boldsymbol{\gamma} \in \Gamma} |F_{\omega'}(\boldsymbol{\gamma}, \Delta_1, n) - \nu| d\boldsymbol{\gamma} d\Delta_1 dn$$

In practice we instead minimize the mean error over a finite training set [23] of points $\mathcal{T} \subset \Gamma \times P$ at which we have pre-computed n and ν using the simulator \mathcal{S} , i.e.,

$$\omega = \operatorname{argmin} \omega' \frac{1}{|\mathcal{T}|} \sum_{(\boldsymbol{\gamma}, \Delta_1, n, \nu) \in \mathcal{T}} |F_{\omega'}(\boldsymbol{\gamma}, \Delta_1, n) - \nu|. \quad (1)$$

To solve this optimization problem, the weights ω are initialized using the heuristic of Ref. [29], then iteratively updated using the first-order stochastic-gradient optimizer ADAM [30] with a minibatch size of 512 and learning rate (step size) of 10^{-3} . We use a DNN with five fully connected layers of 512 units each, with ELU nonlinearities [31], layer normalization [32], and residual connections [33]. For the input layer, we use Fourier feature embedding with four octaves [34]; for the output layer, we use a single linear unit, see Appendix A for details. We select these architectural parameters, and determine when to stop training the DNN, based on its performance on a separate validation set, that is disjoint from \mathcal{T} . We implemented the DNN using the TensorFlow library [35]. Codes for generating data, training, and evaluating the DNN are publicly available [36].

B. Inverse problem

Suppose we have an experimental dataset that provides measurements of penetration capacitance $C_p(D, n)$, which is a quantity that is sensitive to the DOS, as we discuss below. The measurements are acquired while varying (D, n) over some finite set Q ; here D is the strength of an externally applied electric field that affects Δ_1 . The inverse problem is to determine the band structure parameters $\boldsymbol{\gamma}$ for the system, i.e., to find the setting for $\boldsymbol{\gamma}$ for which the simulated results are the closest to the experimental ones. We use the trained DNN F_ω in place of the simulator \mathcal{S} ; we therefore seek

$$\boldsymbol{\gamma} = \operatorname{argmin}_{\boldsymbol{\gamma} \in \Gamma} \min_{\alpha, \beta} \frac{1}{|Q|} \sum_{(D, n) \in Q} d(C_p^*(D, n), C_p(D, n)), \quad (2)$$

where d is a metric quantifying the difference between experimental and predicted values, and we introduced a function $C_p^*(D, n)$ that relies on the electrostatic model to convert the

DOS approximated by the DNN into the penetration field capacitance value,

$$C_p^*(D, n) = \frac{\beta_1}{\beta_2 + F_\omega(\boldsymbol{\gamma}', \alpha_1 D + \alpha_3 D^3, n)}, \quad (3)$$

with $\alpha_{1,3}$ and $\beta_{1,2}$ being parameters that encode the screening and electrostatic characteristics of the experimental setup, respectively. This function and intuition behind these additional parameters is described in detail in Sec. III A.

We solve the minimization problem (2) using the ADAM optimizer [30].² This optimizer makes use of the gradient of the objective with respect to $\boldsymbol{\gamma}$; to calculate this, we use reverse-mode automatic differentiation on the objective in Eq. (2), similar to the back-propagation process used when training the DNN [37], but now with the weights ω held fixed and gradients instead propagated to the inputs $\boldsymbol{\gamma}$.

III. APPLICATION TO ABA GRAPHENE

In this section we discuss a specific application of our method—determining the band structure of Bernal stacked trilayer (ABA) graphene. We first describe the physical system and how it is simulated in Sec. III A. Then, we discuss the dataset that we generate using this simulator (Sec. III B) and use to train the DNN. Finally, we discuss the performance of the trained DNN when solving the forward (Sec. III C) and inverse (Sec. III D) problems, and how we apply it to experimental data (Sec. III E).

A. Physical system and simulation

The band structure of ABA graphene can be decomposed through rotation of the basis into monolayer- and bilayer-graphene-type sectors, which get coupled through a displacement field applied between the layers. The Hamiltonian matrix takes the form [28]

$$H_{\mathbf{k}} = \begin{pmatrix} H_{\text{SLG}} & V \\ V^T & H_{\text{BLG}} \end{pmatrix}, \quad (4)$$

where

$$H_{\text{SLG}} = \begin{pmatrix} \Delta_2 - \frac{\gamma_2}{2} & v_0 \pi^\dagger \\ v_0 \pi & -\frac{\gamma_5}{2} + \delta + \Delta_2 \end{pmatrix},$$

$$H_{\text{BLG}} = \begin{pmatrix} \frac{\gamma_2}{2} + \Delta_2 & \sqrt{2} v_3 \pi & -\sqrt{2} v_4 \pi^\dagger & v_0 \pi^\dagger \\ \sqrt{2} v_3 \pi^\dagger & -2\Delta_2 & v_0 \pi & -\sqrt{2} v_4 \pi \\ -\sqrt{2} v_4 \pi & v_0 \pi^\dagger & \delta - 2\Delta_2 & \sqrt{2} \gamma_1 \\ v_0 \pi & -\sqrt{2} v_4 \pi^\dagger & \sqrt{2} \gamma_1 & \frac{\gamma_5}{2} + \delta + \Delta_2 \end{pmatrix},$$

$$V = \begin{pmatrix} \Delta_1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \Delta_1 \end{pmatrix}.$$

²We also experimented with the constrained trust-region method of Ref. [49], and the quasi-Newton method L-BFGS-B [50], both as implemented in SCIPY [51]. However, noise in the experimental data caused both these methods to perform poorly, becoming trapped in local optima.

Here $v_i = \sqrt{3} a \gamma_i / 2$ and $a = 2.46 \text{ \AA}$ is the lattice constant. γ_i and δ are the hopping and on-site potential parameters of the physical system. $\pi = \xi k_x + i k_y$ is the crystal imaginary momentum measured with respect to K^\pm valley points labeled by $\xi = \pm 1$. Δ_1 describes the effect of the external field and Δ_2 charge asymmetry between external and internal layers.

For producing simulated data we calculate density and DOS on a grid using eigenvalues of Hamiltonian (4) $\epsilon_{\mathbf{k}}$ for each value of chemical potential μ

$$n(\mu) = 4 \frac{S_k}{(2\pi)^2} \frac{1}{N} \sum_{\mathbf{k}} n_F(\epsilon_{\mathbf{k}} - \mu), \quad (5)$$

$$v(\mu) = 4 \frac{S_k}{(2\pi)^2} \frac{1}{N} \sum_{\mathbf{k}} n'_F(\epsilon_{\mathbf{k}} - \mu), \quad (6)$$

where 4 accounts for spin and valley degeneracy, S_k is the area of the Brillouin zone sampled in the grid, N is the number of grid points, and $n_F(x) = 1/(e^{x/T} + 1)$ is the Fermi-Dirac distribution. We use finite temperature $T = 0.025 \text{ meV}$ in the calculation to smoothen singularities of DOS at Van Hove singularity points, and a hexagonal grid with a momentum cutoff $ka = 0.15$ and 7.5×10^5 grid points.

To obtain the dependence of the experimentally measured penetration field capacitance on the DOS we can use the following formula [11,38,39]:

$$C_p = \frac{c_t c_b}{c_t + c_b + e^2 v}, \quad (7)$$

where c_b and c_t is the geometric capacitance of bottom and top gate, e is the charge of the electron. However, this is a simplified formula, which ignores layer polarization change due to the applied electric field D , and views the trilayer graphene as a single-layer system. Therefore, to allow a more general relationship, we introduced the parameters β_1 and β_2 in (3). Besides linear screening, we allow for nonlinear component, so

$$\Delta_1 = \alpha_1 D + \alpha_3 D^3. \quad (8)$$

Due to the presence of a parasitic capacitance, we model the relation between experimentally measured and calculated charge density by

$$n' = n + (\mu - \mu(\Delta_1 = 0, n = 0))/\eta, \quad (9)$$

where η is the inverse of parasitic capacitance.

In the figures, simulation results are presented in terms of inverse of DOS v^{-1} in units of $\text{eV} \cdot A_{\text{u.c.}}$, where $A_{\text{u.c.}} = \sqrt{3} a^2 / 2$ is the area of graphene unit cell. Charge density is always presented in units of cm^{-2} . Band structure parameters γ_i and parameters δ , Δ_1 , Δ_2 , and μ are presented in meV.

B. Data generation

Using the simulation \mathcal{S} described above, we generate data to train and evaluate the DNN. In order to define a suitable grid, we choose the space Γ of valid physical parameters using physical insights into the meaning of the tight-binding parameters. Since δ and γ_5 introduce only a simple energy shift on the spectrum for a considerable range of values, we fix them to $\delta = 35.5 \text{ meV}$ and $\gamma_5 = 50 \text{ meV}$. We also fix $\gamma_0 = 3100 \text{ meV}$ and $\gamma_1 = 380 \text{ meV}$; these are intralayer and direct interlayer hopping matrix elements, which were studied extensively in graphite, and in monolayer and bilayer

TABLE I. Range of physical parameter and grid step used in the calculations.

Parameter	Minimum (meV)	Maximum (meV)	Step (meV)
γ_0	3100	3100	–
γ_1	380	380	–
γ_2	–25	–1	3
γ_3	–340	–242	7
γ_4	100	177	7
γ_5	50	50	–
δ	35.5	35.5	–
Δ_2	0	11	1

graphene. The remaining tuning parameters that constitute $\boldsymbol{\gamma}$ are $\gamma_2, \gamma_3, \gamma_4, \Delta_2$, and η . These parameters are allowed to take values in the ranges shown in Table I. Finally, we define P , the region of acceptable parameters Δ_1 and n , according to $0 \leq \Delta_1 \leq 75$ meV and $|n| < 4 \times 10^{12}$ cm $^{-2}$, corresponding to the physically achievable range of transverse electric fields and densities, respectively.

To generate datasets for training, validation, and testing of the DNN, we iterate over a grid of $\gamma_2, \gamma_3, \gamma_4$, and Δ_2 , with the respective increments for each specified in the last column of Table I. The increment is chosen so that for each parameter the grid contains on the order of 10 points. We also add an additional random offset of up to half the step size to each value, to avoid the DNN overfitting on points of the grid. We then partition this set of parameter vectors into three disjoint subsets: We randomly assign 75% of parameter values to the training set \mathcal{T} , 10% for validation, and the remainder for testing.³

For each value of the structure parameters, we vary Δ_1 and μ in the intervals $[0, 75]$ meV and $[-115, 115]$ meV with steps 2.5 meV and 0.025 meV, respectively. We calculate the corresponding values of n and ν using \mathcal{S} as specified in Eqs. (5)–(6). Since the charge density n for different values of the parasitic capacitance may be determined straightforwardly from that when it is zero, we fix the inverse parasitic capacitance $\eta = \infty$. This yields a dataset of tuples of values $(\gamma_2, \gamma_3, \gamma_4, \Delta_2, \Delta_1, n, \nu)$. We discard tuples where $n \notin P$ (which is beyond our range of interest), or both $|n| < 0.01 \times 10^{12}$ cm $^{-2}$ and $\nu < 0.01$ eV $\cdot A_{u.c}$ (i.e., very close to the charge neutrality point and the gapped region, respectively). During the training process itself, we sample the value of inverse parasitic capacitance, η , from a log-uniform distribution on the interval $[50, 5000] 10^{-12}$ meV \times cm 2 , independently for each element of each minibatch, and transform n according to (9). In total our training set contains 1.38×10^9 points from 14580 settings of the band structure parameters; there are 184×10^6 points in the validation set, and 276×10^6 in the test set.

C. Forward problem: Calculating ν

Using the dataset of simulation results for ABA graphene described above, we train a DNN following the protocol

³Note that the validation set is used for determining convergence and selecting the network architecture, while all results are reported on the testing set; this avoids overfitting to the testing data.

in Sec. II A, to predict the density of states ν . To evaluate the accuracy of the trained DNN, we use it to predict ν for points in the held-out test set, i.e., for parameters $(\gamma_2, \gamma_3, \gamma_4, \Delta_2) \notin \mathcal{T}$ that were not contained in the training dataset. Quantitatively, our method achieves a mean absolute error in ν of 0.00288 eV $\cdot A_{u.c}$, corresponding to a mean relative error of 5.95%. We further analyzed the cause of these errors, by resimulating a random subset 300 of plots in the test set using a finer momentum grid with 6.75×10^6 points and cutoff $ka = 0.2$. Comparing the DNN's predictions to these more accurate simulations, the absolute error reduces to 0.00067 eV $\cdot A_{u.c}$, corresponding to a mean relative error of 1.32%. We also measure the relative error of the lower-resolution simulations forming our dataset, with respect to the 300 higher-resolution plots, and find this is 5.97%. Remarkably, the DNN's predictions are therefore slightly more accurate than the lower-resolution simulations used for its training data. We hypothesize that this is due to the DNN's inductive bias towards learning smooth functions discouraging it from learning the very variable high-frequency artifacts that arise due to the coarser momentum grid. Therefore, in the following figures we show simulator results for higher-resolution momentum grid, although DNN training was performed using lower-resolution data.

Examples of outputs from the DNN, alongside corresponding outputs from the simulator, are shown in Fig. 2. We also visualize the mean absolute difference between the DNN's predictions, and those of the simulator. We see that the DNN produces outputs closely matching the simulator, most importantly accurately incorporating features such as discontinuities and (smoothened) singularities that originate from the changes of Fermi surface topology upon tuning density, n or Δ_1 . For instance, the diagonal feature for negative densities in Fig. 2(a) corresponds to the disappearance of the hole pocket upon increasing density, while parabola-shaped blue line in the negative density region corresponds to the Van Hove singularity where DOS would show a logarithmic divergence in absence of cutoffs introduced due to finite temperature and finite grid size in momentum space. The plots of absolute differences show that the errors are typically very small in regions of uniform ν^{-1} , and are dominated by small misalignments of the features that correspond to DOS discontinuities or singularities.

We also measure the difference in computation time for the simulator and the DNN. Running with four threads on a 3.6 GHz Intel Core i7-9700K CPU, the simulator takes 1445 s to evaluate ν for the grid of (Δ_1, μ) specified in Sec. III B and the fine momentum grid. Running the DNN on the same CPU with four threads takes just 1.1 s. Moreover, it does so with higher-resolution sampling of the range of values of n that are of interest, due to taking n as an input instead of μ , hence avoiding the need to discard calculations where n is outside the range of interest.

D. Inverse problem: Determining parameters

We next consider the inverse problem, of finding the band structure parameters corresponding to a set of measurements. We first evaluate our method on simulated data; this allows validating the proposed approach quantitatively, by measuring

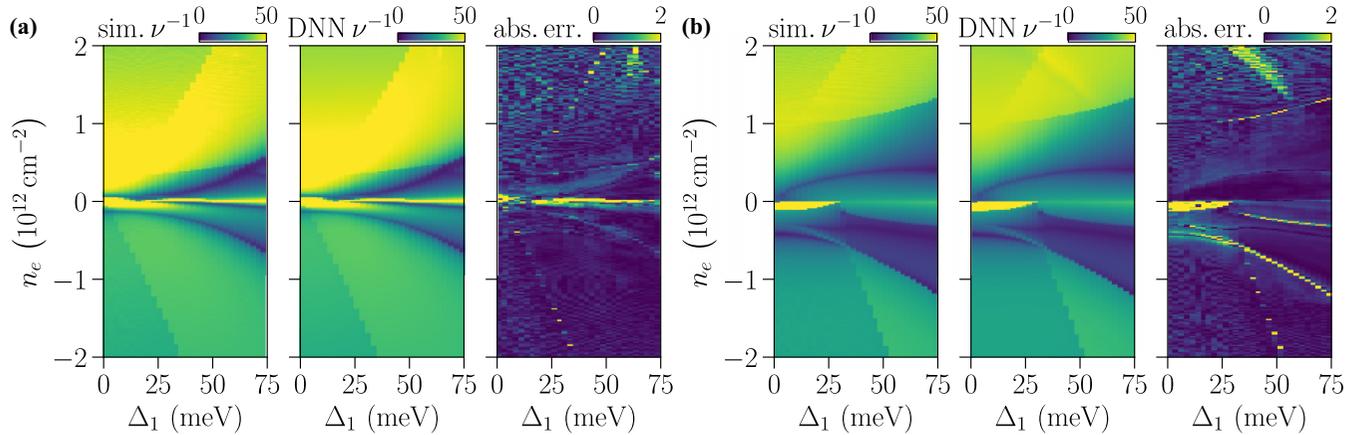


FIG. 2. Predictions of the inverse DOS, ν^{-1} , by our DNN, compared with results from a numerical simulator. (a) and (b) corresponds to the band structure parameters $\gamma_2 = -4.3$, $\gamma_3 = -342$, $\gamma_4 = 133.4$, $\Delta_2 = -0.4$, and $\gamma_2 = -15.6$, $\gamma_3 = -248.8$, $\gamma_4 = 104.9$, $\Delta_2 = 11.2$, respectively. Both of the sets were not present in the training data for the DNN. For each we show: (left) ν^{-1} as a function of field strength Δ_1 and electron charge density n , as predicted by a simulator; (middle) ν^{-1} as predicted by our DNN; ideally this would exactly match the simulator output; (right) the absolute difference between the simulator and DNN values. We see good correspondence between the DNN's outputs and the simulator, indicating that it has been successfully trained to replicate the simulator output, at much lower computational cost.

how accurately it recovers the parameters that were input to the simulator. Specifically, we select 100 simulator-generated plots at random from the held-out test set to use as input; each plot shows how DOS ν varies with Δ_1 and n , and we aim to find the corresponding band structure parameters.

To find the parameters $\boldsymbol{\gamma}$, we follow the procedure described in Sec. II B, using the same DNN as discussed in Sec. III C. As the inputs are simulated data, the DOS itself is directly available, hence we choose the metric d in (2) to be the mean-squared difference between the input and predicted ν .

Table II shows quantitative results from our approach, calculated over 100 plots from the test set. For each band structure parameter, we give the mean and median absolute error, and also the median relative error after first subtracting the midpoint of the corresponding range in Table I. We see that both absolute values of error and also its relative value are small, with the greatest relative error being 5% for the parameter γ_3 , and even smaller for the remaining parameters. Such a confidence interval in determining tight-binding parameters is much smaller compared to the error bars typically available in the literature. For instance, Ref. [11] determines γ_3 parameter with the precision of 17%. Further discussion on quantitative performance of the model using parity plots is provided in Appendix B. Qualitative illustrations of our results for two particular points in the parameter space are given in Fig. 3. We show the plot provided as input to our method, and the result

TABLE II. Quantitative performance on recovering band structure parameters from simulated data.

Parameter	Mean absolute error (meV)	Median absolute error (meV)	Median relative error
γ_2	0.12	0.10	0.0156
γ_3	1.54	1.30	0.0512
γ_4	0.59	0.45	0.0240
Δ_2	0.04	0.03	0.0135

of running the simulator on the predicted parameter values. Here we show high-resolution momentum grid images for both the input and the prediction, although the DNN operated on low-resolution momentum grid images. We see that the residual error is very small, and is less concentrated in the vicinity of DOS discontinuities or singularities, compared to Fig. 2 that benchmarked the forward approach.

E. Application to experimental data

Finally, we apply our method to the high-quality experimental dataset of penetration capacitance measurements for ABA graphene used in Ref. [11]. Note, that Ref. [11] utilized additional datasets, namely, the density of states in presence of magnetic field (Landau level regime) in order to assist in determination of tight-binding parameters and to infer the associated error bars. In contrast, here we use solely zero magnetic field penetration field capacitance data. Our goal is to solve a similar inverse problem to the one in Sec. III D, i.e., predicting the band structure parameters for this system, but now given only noisy measurements of C_p .

Using the mean-squared difference between predicted and experimental C_p for the metric d in Eq. (2) yields poor results in this setting. This is because features in the simulated and experimental plots (e.g., jumps in DOS) do not perfectly align for any choice of parameters, and also because there is a small residual difference in the absolute values of C_p between simulated and experimental data even when the plots are optimally aligned. These discrepancies can be at least partially attributed to multiple physical mechanisms that are not incorporated in our model. In particular, we take account of the screening only phenomenologically by using the third-order expansion (8), whereas the use of self-consistent Hartree-Fock screening would result in a more complex function. In addition, we ignore potential renormalization of the band structure parameters by the strong applied electric fields, that could make all band structure parameters dependent on n and Δ_1 . Finally, disorder and strain effects are not incorporated in

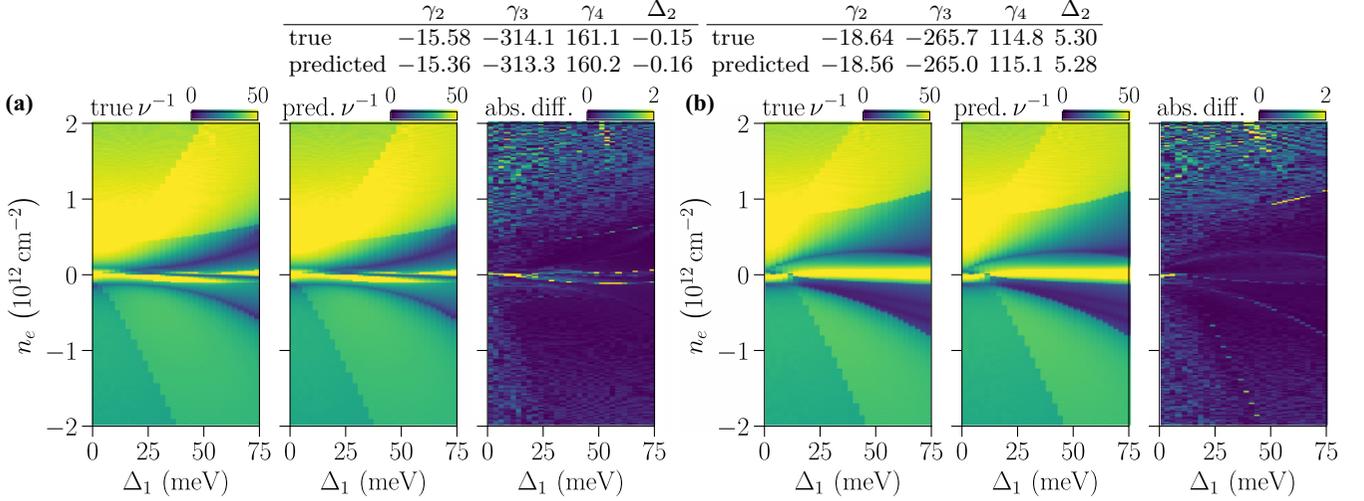


FIG. 3. Results on inverting simulator outputs, i.e., predicting the band structure parameters that gave rise to a given plot. Each panel shows (left) simulated ν^{-1} as a function of Δ_1 and n , for the specified true band parameters; (center) simulated ν^{-1} for the values of the band parameters inferred by our method; (right) absolute difference between simulations and inferred data. The tables at the top compare the true and predicted values of parameters.

our model, yet these may lead to inhomogeneous broadening of the features.

To mitigate this issue, we adopt techniques from computer vision that are used to align dissimilar images; these are typically employed in tasks such as template matching [40,41], medical image registration [42,43], and alignment of satellite imagery [44]. First, in order to match prominent features in the plot such as edges regardless of the absolute values of C_p , we match the derivative of C_p with respect to D and n instead of C_p itself. Second, instead of the mean-squared error, we use the negative zero-normalized cross correlation [41], which

is invariant to changes in mean and standard deviation. We therefore set

$$d(C_p^*, C_p) = - \frac{\partial_D C_p^* - m_D^*}{\sqrt{\langle (\partial_D C_p^* - m_D^*)^2 \rangle}} \cdot \frac{\partial_D C_p - m_D}{\sqrt{\langle (\partial_D C_p - m_D)^2 \rangle}} - \frac{\partial_n C_p^* - m_n^*}{\sqrt{\langle (\partial_n C_p^* - m_n^*)^2 \rangle}} \cdot \frac{\partial_n C_p - m_n}{\sqrt{\langle (\partial_n C_p - m_n)^2 \rangle}} \quad (10)$$

in Eq. (2), where $\langle \cdot \rangle$ is the mean over all (D, n) , $m_D^* = \langle \partial_D C_p^* \rangle$, $m_n^* = \langle \partial_n C_p^* \rangle$, $m_D = \langle \partial_D C_p \rangle$, $m_n = \langle \partial_n C_p \rangle$, and the

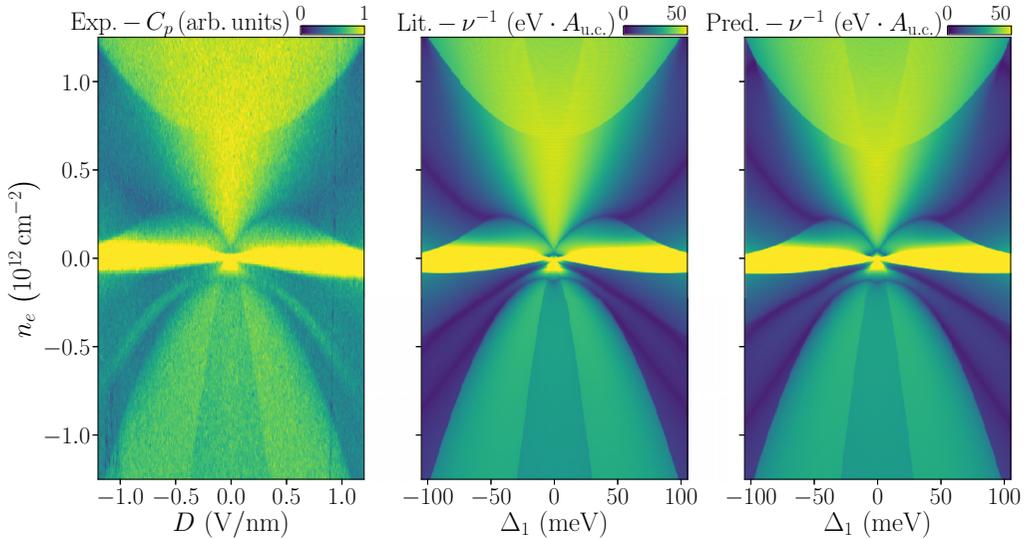


FIG. 4. Inferring band structure parameters from experimental data. The left plot shows experimental C_p measurement as a function of displacement field D and charge density n ; the middle and right plots show the inverse of DOS ν^{-1} as a function of Δ_1 and n for parameters $\gamma_2 = -21$, $\gamma_3 = -290$, $\gamma_4 = 141$, $\Delta_2 = 3.5$ and $\gamma_2 = -15.5$, $\gamma_3 = -312.9$, $\gamma_4 = 132.3$, $\Delta_2 = 3.0$, respectively. The middle figure is the best match to the experimental results cited in the literature [11], while the right one is the prediction of the DNN. The middle and right plots were obtained using simulator with higher resolution of μ ($\Delta\mu = 0.002$), hexagonal grid with maximum momentum $ka = 0.2$, 6.75×10^6 grid points, larger range of Δ_1 and with a step of 0.5.

TABLE III. Quantitative results of sensitivity analysis on experimental data. For each parameter, we record the values to which it must be decreased/increased to cause the distance function $d(C_p^*, C_p)$ to increase by 1%. We specify both absolute and relative values; a large percentage indicates low sensitivity to that parameter/direction, as a large change is required to affect the quality of fit.

Parameter	Decreasing		Best fit (meV)	Increasing	
	meV	rel. diff.		meV	rel. diff.
γ_2	-17.25	-11.2%	-15.51	-14.45	+6.8%
γ_3	-343.5	-9.8%	-312.9	-296.8	+5.1%
γ_4	111.1	-16.1%	132.3	153.2	+15.8%
Δ_2	2.72	-9.5%	3.00	3.17	+5.5%

derivatives of the experimental measurements are approximated using finite differences. Lastly, in order to mitigate against local minima, we average the objective (10) over multiple scales, i.e., subsampling the grid Q by factors 2^i for $i \in \{1 \dots 5\}$, and correspondingly smoothing and down sampling C_p and C_p^* to form scale-space pyramids [45].

Solving the resulting minimization problem as described in Sec. II B recovers the following parameters:

$$\begin{aligned} \gamma_2 &= -15.5 \text{ meV}, & \gamma_3 &= -312.9 \text{ meV}, \\ \gamma_4 &= 132.3 \text{ meV}, & \Delta_2 &= 3.0 \text{ meV}. \end{aligned} \quad (11)$$

These are broadly consistent with the estimates obtained from manual fitting in Ref. [11]. More precisely, our estimates of γ_2 , γ_3 , and γ_4 lie within the error bars of Ref. [11] (although γ_2 is at the upper limit), while Δ_2 is slightly outside (we predict 3.0, compared with their confidence interval [3.3, 3.7]). Notably, the obtained γ_2 is actually very close to the value obtained for ABC trilayer graphene [46].

In Fig. 4, we visualize the experimental data, and the output from the simulator when run with the above parameters at high resolution. The simulated plot shows a good match to the experimental one, with similar features (e.g., discontinuities) appearing at similar locations. We also show the result using the parameters of Ref. [11], these are again qualitatively very similar. From visual inspection the separation between the two Van Hove singularities in the hole region appears to be wider for predicted parameters than for the one from Ref. [11]. This trend is consistent with the experimental results. For all other features the difference between the two models is visually indistinguishable. Therefore, the newly obtained parameters are a better fit to the experiment data.

We also measure the sensitivity of the experimental fit with respect to each of the band structure parameters. Specifically, we increase and decrease each of the parameters γ_2 , γ_3 , γ_4 and Δ_2 separately, until the distance $d(C_p^*, C_p)$ between predicted and experimental C_p values increases by more than 1% from the best fit value. This lets us measure quantitatively how large a change in each parameter is required to cause the same reduction in the quality of fit. Results from this analysis are given in Table III. We see that the fit is least sensitive to γ_4 , with changes of around 16% in either direction required to cause a change of 1% in the distance d . Conversely, increasing γ_3 by just 5.8% leads to a 1% change in d ; Δ_2 and γ_2 are similarly sensitive to increases.

IV. DISCUSSION

To summarize, we proposed a deep neural network-based approach to determining band structure parameters of two-dimensional materials. Our framework consists of two steps, which rely on the existence of a method to simulate the experimentally accessible data. In the first step we train a deep neural network to obtain a more efficient replacement for the data simulator. In the second step, we extract band structure parameters by minimizing the difference between the experimental dataset and the data simulated by the neural network. To illustrate the application of our framework, we focused on a specific material—trilayer graphene—and performed band structure parameter extraction using the experimental data on the penetration field capacitance that effectively probes the density of states. Our procedure resulted in the precise determination of band structure parameters that are close to those determined manually.

In contrast to manual fitting, the approach proposed in our work is more automated, thus involving less human effort. In addition, it is capable of providing more precise values of band structure parameters, and of estimating associated error bars. We expect that our framework can be easily applied with minimal modifications to the penetration field capacitance data in other two-dimensional materials. Moreover, it can be generalized to other experimental probes, such as penetration field capacitance in presence of magnetic field, transport, scanning tunneling microscopy, and other probes that are sensitive to the band structure details. Application to different experimental quantities requires replacement of the simulator function in our framework, which would be straightforward. A more intricate step consists in understanding the relation between the simulated quantities and experimental measurements. For instance, in the case of penetration field capacitance, the matching of experimental data additionally required incorporating nontrivial screening of electric field, the presence of parasitic capacitance and geometric capacitances of the gates. These steps require physical insight into the details of the experimental measurements.

Finally, a conceptually novel generalization direction may include incorporating interaction effects into our framework. For instance, the recent observation of Stoner transitions in twisted bilayer graphene [47,48] or chirally stacked ABC trilayer graphene [46], calls for unambiguous identification of interaction parameters alongside the band structure details (that may be in turn quite significantly renormalized due to interactions). We expect that such an extension of our framework may be feasible, provided one is able to construct an efficient simulator of the density of states that incorporates the interactions. This would constitute an important step towards an experimental extraction of the complete Hamiltonian governing electronic degrees of freedom, thus bringing two-dimensional materials on par with quantum simulators that use cold atoms or other platforms to synthetically engineer and verify a desired Hamiltonian.

ACKNOWLEDGMENTS

A.F.Y. acknowledges primary support from the Department of Energy under Award No. DE-SC0020043, and additional

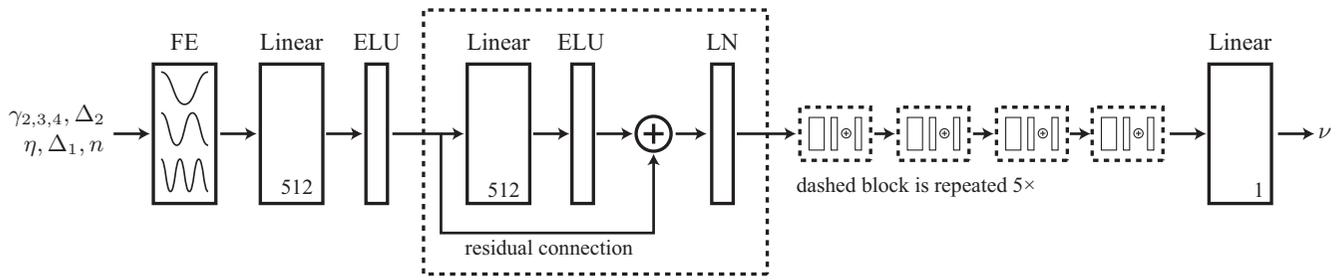


FIG. 5. Architecture of the deep neural network F_ω . The DNN takes band structure parameters ($\gamma_2, \gamma_3, \gamma_4, \Delta_2$), inverse parasitic capacitance η , interlayer asymmetry Δ_1 and charge density n as input. It processes them with a Fourier embedding layer (FE), linear layer with 512 outputs, and ELU nonlinearity, followed by five blocks with residual connections consisting of linear, ELU, and layer-norm (LN) layers. The resulting features are mapped via a linear projection to the density of states ν .

support from the Gordon and Betty Moore Foundation under Award No. GBMF9471 for group operations.

APPENDIX A: ADDITIONAL MODEL AND TRAINING DETAILS

The architecture of our deep neural network F_ω is shown in Fig. 5. In total, the DNN has 1.35×10^6 weights (i.e., trainable parameters). We train the model using ADAM [30] with a fixed learning rate of 10^{-3} , and the standard momentum parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-7}$. The batch size is set to 512. Convergence is determined according to the loss on the held-out validation set—we stop training once the validation loss has failed to decrease for ten consecutive epochs, and retain the checkpoint with minimum validation loss. We use a single Nvidia RTX 2080 Ti GPU for training, operating at 32-bit precision.

APPENDIX B: ADDITIONAL QUANTITATIVE RESULTS ON THE INVERSE PROBLEM

To gain further insights on the performance of the model for the inverse problem, Fig. 6 shows parity plots of the band structure parameters $\gamma_2, \gamma_3, \gamma_4$, and Δ_2 for the same test set of 100 plots, which was used for obtaining the metrics in Table II. As noted in the discussion of Table II the largest deviation is observed for parameter γ_3 , and Fig. 6(b) shows that the prediction of the model is always higher than the true values. This systematic error hints at a direction for further

improving the model accuracy, and would be an interesting target for future investigation.

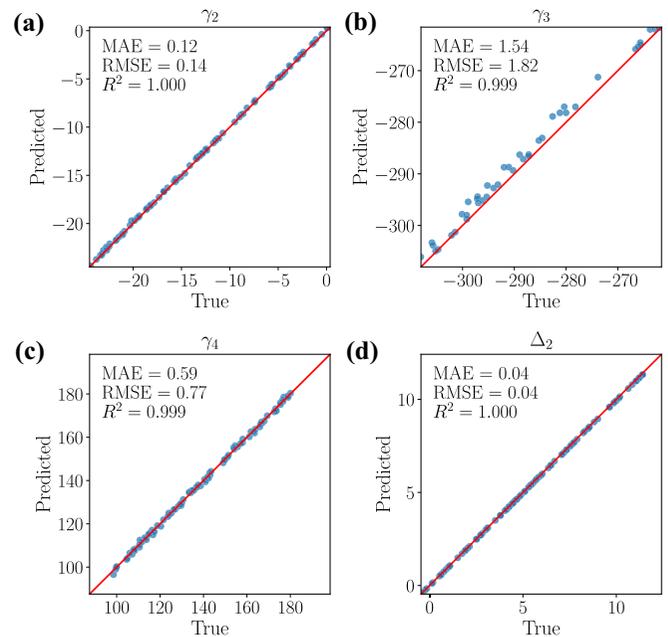


FIG. 6. Parity plots of the band structure parameters from simulated data for the inverse problem using the test set of 100 plots. The analyzed data is the same as in Table II. MAE (mean absolute error), RMSE (root mean-square error) and R^2 metrics are also shown. Equality line ($y = x$) is shown in red.

- [1] W. A. Harrison, *Electronic Structure and the Properties of Solids: The Physics of the Chemical Bond* (Dover Publications, New York, 1989).
- [2] P. J. Hasnip, K. Refson, M. I. Probert, J. R. Yates, S. J. Clark, and C. J. Pickard, Density functional theory in the solid state, *Philos. Trans. R. Soc. Lond. A* **372**, 20130270 (2014).
- [3] A. Damascelli, Z. Hussain, and Z.-X. Shen, Angle-resolved photoemission studies of the cuprate superconductors, *Rev. Mod. Phys.* **75**, 473 (2003).
- [4] F. de Groot, High-resolution x-ray emission and x-ray absorption spectroscopy, *Chem. Rev.* **101**, 1779 (2001).
- [5] D. Shoenberg, *Magnetic Oscillations in Metals* (Cambridge University Press, Cambridge, 2009).
- [6] F. Wooten, *Optical Properties of Solids* (Academic Press, New York, 2013).
- [7] U. Mizutani, *Introduction to the Electron Theory of Metals* (Cambridge University Press, Cambridge, 2001).
- [8] K. Novoselov, A. Mishchenko, A. Carvalho, and A. Castro Neto, 2D materials and van der Waals heterostructures, *Science* **353**, aac9439 (2016).
- [9] J. P. Eisenstein, L. N. Pfeiffer, and K. W. West, Negative Compressibility of Interacting Two-Dimensional Electron and Quasiparticle Gases, *Phys. Rev. Lett.* **68**, 674 (1992).
- [10] J. Island, X. Cui, C. Lewandowski, J. Khoo, E. Spanton, H. Zhou, D. Rhodes, J. Hone, T. Taniguchi, K. Watanabe, L. S. Levitov, M. P. Zaletel, and A. F. Young, Spin-orbit-driven band inversion in bilayer graphene by the van der Waals proximity effect, *Nature (London)* **571**, 85 (2019).
- [11] A. A. Zibrov, P. Rao, C. Kometter, E. M. Spanton, J. I. A. Li, C. R. Dean, T. Taniguchi, K. Watanabe, M. Serbyn, and A. F. Young, Emergent Dirac Gullies and Gully-Symmetry-Breaking Quantum Hall States in ABA Trilayer Graphene, *Phys. Rev. Lett.* **121**, 167601 (2018).
- [12] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [13] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [14] K. Choo, G. Carleo, N. Regnault, and T. Neupert, Symmetries and Many-Body Excitations with Neural-Network Quantum States, *Phys. Rev. Lett.* **121**, 167204 (2018).
- [15] Z. Cai and J. Liu, Approximating quantum many-body wave functions using artificial neural networks, *Phys. Rev. B* **97**, 035116 (2018).
- [16] A. Borin and D. A. Abanin, Approximating power of machine-learning ansatz for quantum many-body states, *Phys. Rev. B* **101**, 195141 (2020).
- [17] T. Christensen, C. Loh, S. Picek, D. Jakobovic, L. Jing, S. Fisher, V. Ceperic, J. D. Joannopoulos, and M. Soljacic, Predictive and generative machine learning models for photonic crystals, *Nanophotonics* **9**, 4183 (2020).
- [18] A. Nikulin, I. Zisman, M. Eich, A. Y. Petrov, and A. Itin, Machine learning models for photonic crystals band diagram prediction and gap optimisation, *Photon. Nanostruct. Fund. Appl.* **52**, 101076 (2022).
- [19] A. Sheverdin, F. Monticone, and C. Valagiannopoulos, Photonic Inverse Design with Neural Networks: The Case of Invisibility in the Visible, *Phys. Rev. Appl.* **14**, 024054 (2020).
- [20] B. S. Rem, N. Käming, M. Tarnowski, L. Asteria, N. Fläschner, C. Becker, K. Sengstock, and C. Weitenberg, Identifying quantum phase transitions using artificial neural networks on experimental data, *Nat. Phys.* **15**, 917 (2019).
- [21] Y. Zhang, A. Mesaros, K. Fujita, S. Edkins, M. Hamidian, K. Ch'ng, H. Eisaki, S. Uchida, J. S. Davis, E. Khatami, and E.-A. Kim, Machine learning in electronic-quantum-matter imaging experiments, *Nature (London)* **570**, 484 (2019).
- [22] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, 2016), <http://www.deeplearningbook.org>.
- [23] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer-Verlag, Berlin, 2006).
- [24] K. Hornik, M. Stinchcombe, and H. White, Multilayer feedforward networks are universal approximators, *Neural Networks* **2**, 359 (1989).
- [25] P. Kidger and T. Lyons, Universal approximation with deep narrow networks, in *Proceedings of Thirty Third Conference on Learning Theory*, edited by J. Abernethy and S. Agarwal (PMLR, 2020), Vol. 125, pp. 2306–2327.
- [26] V. Vapnik, Principles of risk minimization for learning theory, *Adv. Neural Info. Proc. Syst.* **4**, 831 (1991).
- [27] M. Dresselhaus and G. Dresselhaus, Intercalation compounds of graphite, *Adv. Phys.* **30**, 139 (1981).
- [28] M. Serbyn and D. A. Abanin, New Dirac points and multiple Landau level crossings in biased trilayer graphene, *Phys. Rev. B* **87**, 115422 (2013).
- [29] X. Glorot and Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, edited by Y. W. Teh and M. Titterton, *Proceedings of Machine Learning Research*, Vol. 9 (PMLR, Sardinia, Italy, 2010), pp. 249–256.
- [30] D. P. Kingma and J. L. Ba, Adam: A method for stochastic optimization, in *International Conference on Learning Representations* (ICLR, Ithaca, NY, 2015).
- [31] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, Fast and accurate deep network learning by exponential linear units (ELUs), [arXiv:1511.07289](https://arxiv.org/abs/1511.07289).
- [32] J. L. Ba, J. R. Kiros, and G. E. Hinton, Layer normalization, [arXiv:1607.06450](https://arxiv.org/abs/1607.06450).
- [33] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, New York, 2016), pp. 770–778.
- [34] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng, Fourier features let networks learn high frequency functions in low dimensional domains, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Red Hook, NY, 2020), Vol. 33, pp. 7537–7547.
- [35] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems, [arXiv:1603.04467](https://arxiv.org/abs/1603.04467).
- [36] <https://git.ista.ac.at/qdyn/ML-band-structure-DOS>.

- [37] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning representations by back-propagating errors, *Nature (London)* **323**, 533 (1986).
- [38] B. M. Hunt, J. I. A. Li, A. A. Zibrov, L. Wang, T. Taniguchi, K. Watanabe, J. Hone, C. R. Dean, M. Zaletel, R. C. Ashoori, and A. F. Young, Direct measurement of discrete valley and orbital quantum numbers in bilayer graphene, *Nature Commun.* **8**, 948 (2017).
- [39] A. A. Zibrov, C. Kometter, H. Zhou, E. Spanton, T. Taniguchi, K. Watanabe, M. Zaletel, and A. Young, Tunable interacting composite fermion phases in a half-filled bilayer-graphene Landau level, *Nature (London)* **549**, 360 (2017).
- [40] R. M. Haralick and L. G. Shapiro, *Computer and Robot Vision, Volume II* (Addison-Wesley, Boston, 1992), pp. 316–317.
- [41] J. P. Lewis, Fast Template Matching, *Vision Interface 95* (Canadian Image Processing and Pattern Recognition Society, Quebec City, Canada, 1995), pp. 120–123.
- [42] C. Studholme, D. L. G. Hill, and D. J. Hawkes, Automated 3-D registration of MR and CT images of the head, *Medical Image Analysis* **1**, 163 (1996).
- [43] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. Pluim, Elastix: A toolbox for intensity-based medical image registration, *IEEE Trans. Med. Imag.* **29**, 196 (2010).
- [44] L. G. Brown, A survey of image registration techniques, *ACM Comput. Surv.* **24**, 325 (1992).
- [45] C. H. Anderson, J. R. Bergen, P. J. Burt, and J. M. Ogden, Pyramid methods in image processing, *RCA Engineer* **29**, 33 (1984).
- [46] H. Zhou, T. Xie, A. Ghazaryan, T. Holder, J. R. Ehrets, E. M. Spanton, T. Taniguchi, K. Watanabe, E. Berg, M. Serbyn, and A. F. Young, Half-and quarter-metals in rhombohedral trilayer graphene, *Nature (London)* **598**, 429 (2021).
- [47] U. Zondiner, A. Rozen, D. Rodan-Legrain, Y. Cao, R. Queiroz, T. Taniguchi, K. Watanabe, Y. Oreg, F. von Oppen, A. Stern, E. Berg, P. Jarillo-Herrero, and S. Ilani, Cascade of phase transitions and Dirac revivals in magic-angle graphene, *Nature (London)* **582**, 203 (2020).
- [48] S. Wu, Z. Zhang, K. Watanabe, T. Taniguchi, and E. Y. Andrei, Chern insulators, van Hove singularities and topological flat bands in magic-angle twisted bilayer graphene, *Nat. Mater.* **20**, 488 (2021).
- [49] R. H. Byrd, M. E. Hribar, and J. Nocedal, An interior point algorithm for large-scale nonlinear programming, *SIAM J. Optim.* **9**, 877 (1999).
- [50] C. Zhu, R. H. Byrd, and J. Nocedal, Algorithm 778: L-BFGS-B: FORTRAN routines for large scale bound constrained optimization, *ACM Trans. Math. Software* **23**, 550 (1997).
- [51] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).