





Deep-learning density functionals for gradient descent optimizationE. Costa ^{1,2} G. Scriva ^{1,2} R. Fazio ^{3,4} and S. Pilati ^{1,2}¹*School of Science and Technology, Physics Division, Università di Camerino, 62032 Camerino, Italy*²*INFN-Sezione di Perugia, 06123 Perugia, Italy*³*Abdus Salam ICTP, Strada Costiera 11, I-34151 Trieste, Italy*⁴*Dipartimento di Fisica, Università di Napoli “Federico II,” Monte S. Angelo, I-80126 Napoli, Italy*

(Received 23 May 2022; accepted 27 September 2022; published 21 October 2022)

Machine-learned regression models represent a promising tool to implement accurate and computationally affordable energy-density functionals to solve quantum many-body problems via density functional theory. However, while they can easily be trained to accurately map ground-state density profiles to the corresponding energies, their functional derivatives often turn out to be too noisy, leading to instabilities in self-consistent iterations and in gradient-based searches of the ground-state density profile. We investigate how these instabilities occur when standard deep neural networks are adopted as regression models, and we show how to avoid them by using an ad hoc convolutional architecture featuring an interchannel averaging layer. The main testbed we consider is a realistic model for noninteracting atoms in optical speckle disorder. With the interchannel average, accurate and systematically improvable ground-state energies and density profiles are obtained via gradient-descent optimization, without instabilities nor violations of the variational principle.

DOI: [10.1103/PhysRevE.106.045309](https://doi.org/10.1103/PhysRevE.106.045309)**I. INTRODUCTION**

Density functional theory (DFT) is the workhorse of computational material science and quantum chemistry [1]. It is based on rigorous theorems [2–4] certifying that the ground-state energy can be computed minimizing a (generally unknown) functional of the density profile, allowing us to bypass computationally prohibitive wave-function-based methods. However, the available approximations for the density functional are reliable only for weakly correlated materials, while in the regime of strong electron correlations dramatic failures may occur [5]. In recent years, machine-learning (ML) algorithms have reached remarkable breakthroughs in various branches of physics research [6–9], and they have also been adopted in the framework of DFT, both for continuous-space [10–14] and for one-dimensional tight-binding models [15–17]. These algorithms pave the way to databased approaches to the development of density functionals. Furthermore, they facilitate the implementation of computationally convenient strategies based on orbital-free DFT [18,19]. Most previous studies adopted relatively simple regression models, such as, e.g., kernel ridge regression, showing that moderately large training sets of ground-state density profiles and corresponding energies allow us to reconstruct remarkably accurate density functionals. Also artificial neural networks have been adopted (see, e.g., [13,19]), considering standard architectures such as the convolutional neural networks (CNNs) popular in the field of computer vision. Unfortunately, in the case of continuous-space models, severe drawbacks have emerged when such ML functionals have been employed in self-consistent calculations and in gradient-based optimizations. Specifically, the functional derivatives

turned out to be too noisy, leading to unphysical density profiles and to strong violation of the variational principle [10,19,20]. Some remedial strategies have already been explored. Essentially, they resort to gradient denoising via dimensionality reduction [20] or basis truncation [21], to constrained optimization, or they aim at exploiting additional information (e.g., energy derivatives) in the training process [14,19]. These strategies have provided significant benefits, but they have some limitations, as they might lead to variational biases or they require additional data that are far less accessible. Due to the pivotal role played by DFT, further complementary strategies are highly desirable.

In this article, we investigate the use of deep neural networks as regression models to reconstruct continuous-space density functionals from training data. Our main finding is that a tailored convolutional network featuring an interchannel averaging operation allows us to avoid the drawbacks mentioned above. Following analogous previous studies [10,11,18–20], the testbed we consider is a single-particle model, but we mostly focus on a more realistic Hamiltonian that describes ultracold atoms moving in one-dimensional optical speckle patterns. Our analysis is complemented by addressing deep-well models from the literature (see the Appendix). Our aim is to develop a sufficiently effective deep-learned functional to allow searching the ground-state energy and the density profile of previously unseen instances of speckle disorder via gradient-descent optimization. We show that the most popular network architectures, namely, the standard CNNs, are inadequate for this task. Indeed, while they provide remarkably accurate energy predictions when fed with exact ground-state density profiles, their functional derivatives are too noisy. This leads to instabilities in the gradient-descent search of

the density profile which minimize the energy functional, unless the accuracy is jeopardized via an early halting of the optimization procedure. We demonstrate that these instabilities can be avoided with the tailored neural network. This is inspired by an ensemble-averaging mechanism, and it features, beyond the standard multichannel convolutional layers, additional layers that perform an interchannel averaging operation. We show that this feature allows us to iterate gradient-descent steps at will, providing accurate results that can be systematically improved by increasing the number of channels.

The rest of the article is organized as follows: In Sec. II we describe the formalism of DFT based on deep learning, the structure of the standard and of the average-channel neural networks, as well as the gradient-descent technique used to find ground-state energies and densities. The main testbed model we address is described in Sec. III. Therein we also report details on the dataset, on the protocol used for network training, and on the accuracy reached in the regression task. The results obtained in gradient-descent optimization with the standard and with the average-channel neural networks are compared in Sec. IV. The instabilities occurring with standard networks are highlighted, and their suppression with the inclusion of the average layer is discussed in detail. Section V provides a summary of the main findings and some comments on future perspectives. To favor comparison with previous studies [10,18], in Appendix we report the test of the average-channel neural network in deep potential wells defined by the sum of three Gaussian functions.

II. DENSITY FUNCTIONAL THEORY WITH ARTIFICIAL NEURAL NETWORKS

ML provides novel promising approaches to learn energy-density functionals for DFT from data. These functionals have the potential to accurately describe strongly correlated systems. However, their variational minimization to search for the ground-state density profile turned out to be problematic due to noisy functional derivatives [10,11,20]. This problem is already evident in noninteracting systems. Henceforth, in the following we focus on single-particle problems, but the technique we develop can be applied to interacting systems via the creation of suitable training sets.

In this article, we consider one-dimensional single-particle Hamiltonians written in the form

$$H = -\frac{\hbar^2}{2m} \frac{d^2}{dx^2} + V(x), \quad (1)$$

where \hbar is the reduced Planck constant and m is the particle mass. The external potential $V(x)$ is compatible with the adopted periodic boundary conditions. In the framework of DFT, one aims at computing the ground-state energy e_{gs} of the Hamiltonian (1) from a functional of the density profile: $e_{\text{gs}} = E[n_{\text{gs}}]$. Here, n_{gs} indicates the density profile,

$$n_{\text{gs}}(x) = |\psi_{\text{gs}}(x)|^2, \quad (2)$$

where $\psi_{\text{gs}}(x)$ is the ground-state wave function. The first Hohenberg-Kohn theorem guarantees that, in principle, this

TABLE I. Coefficient of determination R^2 for two standard CNNs with ReLU and with Softplus activation functions, and for three CNNs with different numbers of channels N_c . The test set includes 15 000 (previously unseen) instances of the optical speckle pattern. The first column reports the number of convolutional blocks.

Blocks	Activation	N_c	k_s	Neural network	R^2
2	ReLU	60	[4,2]	$\overline{\text{CNN}}$	0.99961
2	ReLU	140	[4,2]	$\overline{\text{CNN}}$	0.99969
2	ReLU	260	[4,2]	$\overline{\text{CNN}}$	0.99967
3	ReLU	30	2	CNN	0.99996
3	Softplus	30	2	CNN	0.99991

functional exists [2]. In practice, it is convenient to separate the known potential energy contribution, seeking a functional for the kinetic energy only [22]:

$$t_{\text{gs}} = e_{\text{gs}} - \int_0^L dx n_{\text{gs}}(x) V(x) \equiv T[n_{\text{gs}}]. \quad (3)$$

It is worth pointing out that we do not adopt the Kohn-Sham formalism. As in previous studies on ML-based DFT [10,11,18–20], the orbital-free formalism is used, attempting to approximate the kinetic energy (eventually together with energy terms in interacting systems) as a functional of the density. If successful, this attempt would therefore also lead to a significant reduction of computational cost compared to the more demanding Kohn-Sham approach.

Deep-learning techniques can be adopted in the DFT framework. The first task is to train a deep neural network to map ground-state density profiles to the corresponding kinetic energies, therefore learning the unknown functional $T[n_{\text{gs}}]$. This can be achieved via supervised learning from a dataset including many instances of density profiles associated with the corresponding kinetic energies, $\{n_{\text{gs},k}, t_{\text{gs},k}\}$. The integer k labels the instances in the dataset. The parameters of the neural network, collectively denoted with ω , are optimized by minimizing the loss function, namely the mean-squared error

$$\mathcal{L}(\omega) = \frac{1}{N_{\text{train}}} \sum_{k=1}^{M_{\text{train}}} |t_{\text{gs},k} - \tilde{T}_{\text{gs},k}(\omega)[n_{\text{gs},k}]|^2, \quad (4)$$

where $\tilde{T}_{\text{gs},k}(\omega)[n_{\text{gs},k}]$ denotes the kinetic energy predicted by the neural network, and N_{train} is the number of instances in the training set. This optimization can be performed using the stochastic gradient descent algorithms or one of its many successful variants.

A. Neural networks

The first regression model we consider is a standard CNN. Its structure is familiar from many fields where deep learning has proven successful, such as, e.g., image recognition. It is composed of $N_b = 3$ convolutional blocks. Each block includes a convolutional layer with N_c channels (this hyperparameter is specified in Table I), whose filter size is $k_f = 13$ and padding type is periodic, and an average pooling layer with a kernel size $k_s = 2$. Two variants of this network are considered, using two popular activation functions, namely

the ReLU function, defined as

$$\text{ReLU}(x) = \begin{cases} x & \text{if } x > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

and the Softplus function

$$\text{Softplus}(x) = \ln(1 + \exp(x)). \quad (6)$$

The last convolutional block is processed through a flattening layer and then connected to a dense layer with only one neuron (identity activation function) to generate a scalar output. As discussed in detail in Sec. IV, these standard CNNs turn out to be inadequate for the DFT framework. In particular, their functional derivatives are too noisy to perform a gradient-based search of the ground state. Therefore, we introduce a novel tailored architecture that features an inter-channel average operation. In the following, this network will be referred to as average-channel CNN (CNN). Specifically, this model is composed of two convolutional blocks, each including N_c convolutional channels ($k_f = 13$ and periodic padding), an average pooling layer ($k_s = 4$ in the first block and $k_s = 2$ in the second block), and, notably, an additional layer, where each neuron computes the average of the activations of the corresponding neurons in all channels of the previous layer. The activation function is ReLU. The last convolutional block passes a flattening layer and is then connected to one dense layer with only one neuron, as in the standard CNN case. It is worth pointing out that this average operation reduces the scaling of the number of parameters from quadratic to linear in N_c . This allows us to consider architectures with many channels without facing prohibitive computational costs nor overfitting problems. In each architecture adopted in this article, all convolutional blocks feature the same number of channels. Explorations performed with different numbers lead to similar findings, so we do not discuss them to avoid burdening the presentation. Hereafter, we describe the operations performed by the convolutional blocks more formally.

In a standard CNN, the action of the n th convolutional block corresponds to the following convolution operation:

$$h_n^\alpha(x) = \frac{1}{k_s} \int_{x-k_s/2}^{x+k_s/2} dy \times \text{act} \left[\sum_{\beta} \int_{y-k_f/2}^{y+k_f/2} dx' W_n^{\alpha\beta}(y, x') h_{n-1}^\beta(x') + v_n^\alpha \right], \quad (7)$$

where act is the chosen activation function, the matrices $W_n^{\alpha\beta}$ represent, for each filter α and input channel β , a kernel of size k_f , and v_n^α are the set of biases for each filter. Instead, in the CNN, the n th block has an additional interchannel average operation, which is expressed as

$$\bar{h}_n(x) = \frac{1}{N_c} \sum_{\alpha} h_n^\alpha(x), \quad (8)$$

where

$$h_n^\alpha(x) = \frac{1}{k_s} \int_{x-k_s/2}^{x+k_s/2} dy \times \text{act} \left[\int_{y-k_f/2}^{y+k_f/2} dx' W_n^\alpha(y, x') \bar{h}_{n-1}(x') + v_n^\alpha \right] \quad (9)$$

represents the previous N_c parallel convolution operations. Notice that in Eqs. (7) and (8) integrals actually indicate discrete operations. All the neural networks considered in this work are implemented and trained using the Pytorch library, exploiting automatic differentiation to compute discrete functional derivatives [23].

B. Formalism of gradient-descent optimization

Once the kinetic energy functional has been learned by the neural network, and consequently we assume $T[n] \equiv \tilde{T}_\omega[n]$, both the ground-state energy and the density corresponding to a new instance of the Hamiltonian can be obtained from the variational principle. Indeed, the second Hohenberg-Kohn theorem ensures that the (exact) functional is minimized by the ground-state density profile [2]. This can be expressed using the Euler equation:

$$\frac{\delta T[n]}{\delta n(x)} + V(x) - \mu = 0, \quad (10)$$

where μ is a normalization constraint. Its solution can be efficiently obtained using the gradient-descent algorithm, as usually done in orbital-free DFT [24]. Specifically, one iterates the following update rule:

$$n_{t+1}(x) = n_t(x) - \eta \left(\frac{\delta T[n_t]}{\delta n(x)} + V(x) - \mu_t \right), \quad (11)$$

starting from a reasonably chosen initial profile $n_0(x)$. In this equation, $\eta > 0$ is the chosen learning rate, the integer $t = 0, 1, 2, \dots, t_{\max}$ labels the steps, and the adaptive coefficient μ_t is introduced to ensure the normalization condition:

$$\int dx n_t(x) = 1. \quad (12)$$

To ensure the density never becomes negative, it is convenient to perform a variable change in Eq. (11): $\chi = \sqrt{n(x)}$. This leads to the following update rule:

$$\chi_{t+1}(x) = \chi_t(x) - \eta \left(\frac{\delta T[\chi_t^2]}{\delta \chi(x)} + 2\chi_t(x)V(x) - 2\chi_t(x)\mu_t \right), \quad (13)$$

where the normalization coefficient is computed as

$$\mu_t = \frac{\int dx \left(\frac{1}{2} \frac{\delta T[\chi_t^2]}{\delta \chi_t(x)} \chi_t(x) + \chi_t^2(x)V(x) \right)}{\int dx \chi_t^2(x)}. \quad (14)$$

The computation of the square of the function χ is performed by an additional layer in Pytorch. Henceforth, the functional derivative with respect to χ can be directly computed exploiting automatic differentiation.

Whether the gradient descent algorithm reaches the ground-state or not depends on two major issues. First, the

optimization might get stuck in a local minimum. Indeed, the optimization landscape is not proven to be convex, even for the (unknown) exact functional. Convexity can be instead proven for an extended functional, defined in a domain including density profiles not corresponding to ground states [4]. This problem can be mitigated by repeated the minimization process starting from different initial profiles, or by introducing random steps based on, e.g., Metropolis-type algorithms. The second issue is the accuracy of the functional derivative. Noisy and inaccurate derivatives might create unphysical density profiles, clearly not corresponding to ground states. Since the regression model was not trained on such profiles, it might provide very inaccurate energy predictions, even lower than the exact ground-state energy e_{gs} . This leads to dramatic failures of the gradient descent optimization, even to large violations of the variational principle. This problem has already been emphasized in the literature, and it was indicated as a major challenge to be overcome for the further development of ML-based DFTs. In Refs. [10,11,20], the adopted regression model was kernel ridge regression. This model typically requires smaller datasets for training, but it is less efficient than the deep neural networks in systematically extracting further information from larger and larger datasets. Even kernel ridge regression led to noisy derivatives. To circumvent this problem, the authors introduced two main techniques, referred to as local principal component analysis and nonlinear gradient denoising. They aim at projecting the functional derivative to the manifold tangent to the one spanned by ground-state density profiles. Other works included derivative data in the training process—also adopting standard CNNs—using the Sobolev Loss [18,19]. We emphasize that our goal is to train the regression model using only ground-state density profiles and the corresponding energies, avoiding resorting to less accessible data such as excited-state properties or energy gradients.

III. TESTBED MODEL AND TRAINING DATASET

The main testbed model we consider is the single-particle model (1), where the (random) external potential $V(x)$ is designed to represent the effect of optical speckle patterns on ultracold atoms. Notice that another testbed, borrowed from the literature, is considered in Appendix, allowing us to further characterize the domain of applicability of the CNN. The speckle potentials can be created by applying a specific filter in Fourier space to a random complex Gaussian field [27]. The filter corresponds to the aperture of the optical apparatus used to experimentally create the field, and it fixes the characteristic size of the speckle grains. In fact, with this choice the Hamiltonian (1) describes early cold-atom experiments on Anderson localization in one dimension [25,26]. The statistical and spectral properties of optical speckle patterns are known [27–29]. The intensity of the potential V in a point x follows the probability distribution

$$P(V) = \exp\left(-\frac{V}{V_0}\right) \quad (15)$$

for $V \geq 0$, and $P(V) = 0$ otherwise; $V_0 \geq 0$ is the average intensity, and it also coincides with the standard deviation. It is the unique parameter determining the disorder strength.

The two-point autocorrelation function satisfies the following equation:

$$\frac{\langle V(x' + x)V(x') \rangle}{V_0^2} - 1 = \frac{\sin(\pi x/\gamma)^2}{(\pi x/\gamma)^2}, \quad (16)$$

where γ determines the correlation length, namely the size of the typical speckle grains. In the above equation, the brackets $\langle \cdot \rangle$ indicate the average over many random realizations of the speckle pattern. This ensemble average coincides with the spatial average for sufficiently large systems. The correlation energy $E_c = \frac{\hbar^2}{2m\gamma}$ separates the strong disorder regime $V_0 \gg E_c$, where the low-energy orbitals are localized on a lengthscale of order γ [30] due to the Anderson localization phenomenon [31], from the weak disorder regime $V_0 \ll E_c$, where their localization length is much larger. Notice that in one dimension, any disorder strength induces Anderson localization in a sufficiently large system [32]. In the following, we consider the relatively large system size $L = 14\gamma$ and the intermediate disorder strength $V_0 = 0.5E_c$. This choice allows us to generate rather variegated ground-state density profiles with different shapes and varying degrees of localization, depending on the details of the specific realization of the speckle pattern. Therefore, the Hamiltonian (1) represents a stringent testbed for the DFT framework. Two representative instances of the speckle pattern are shown in Fig. 1, together with the corresponding ground-state density profiles.

Different random realizations of the speckle pattern can be efficiently generated on a discrete grid with the algorithm described in Refs. [33,34]. We choose a fine grid with $N_g = 256$ points, such that the grid step $\delta_x = L/N_g \ll \gamma$. The ground-state energy e_{gs} and the corresponding orbital $\psi_{\text{gs}}(x)$ are determined via exact diagonalization using a high-order finite-difference formula. Computations performed with finer grids show that the discretization error is negligible. The choice of such a fine grid allows us to compute all spatial integrals [see, e.g., those in Eqs. (3) and (14)] with the discrete approximation $\int_0^L dx \rightarrow \delta_x \sum_{i=1}^{N_g}$. Higher-order approximations lead to essentially indistinguishable results for our purposes. Furthermore, the functional derivative in Eq. (13) is computed as

$$\frac{\delta T[\psi]}{\delta \psi(x)} = \frac{\partial T[\psi]}{\partial \psi(x_i)} \frac{1}{\delta_x}. \quad (17)$$

For this, we exploit Pytorch automatic differentiation. The training of the neural networks, namely the minimization of the loss function Eq. (4), is performed using the Adam algorithm [35]. The chosen learning rate is $l_r = 10^{-4}$, the minibatch size is $N_b = 100$ and the training epochs are $N_e = 1200$. The other parameters of the Adam algorithm are set at their suggested default values. Our global dataset is composed of 150 000 instances. As is customary in deep-learning studies, we split it into a training set (81%), a validation set (9%), and a test set (10%). To measure the accuracy in the regression task, we consider the coefficient of determination, defined as

$$R^2 = 1 - \frac{\sum_{k=1}^{N_{\text{test}}} |t_{\text{gs},k} - \tilde{T}_\omega[n_{\text{gs},k}]|^2}{N_{\text{test}} \sigma^2}, \quad (18)$$

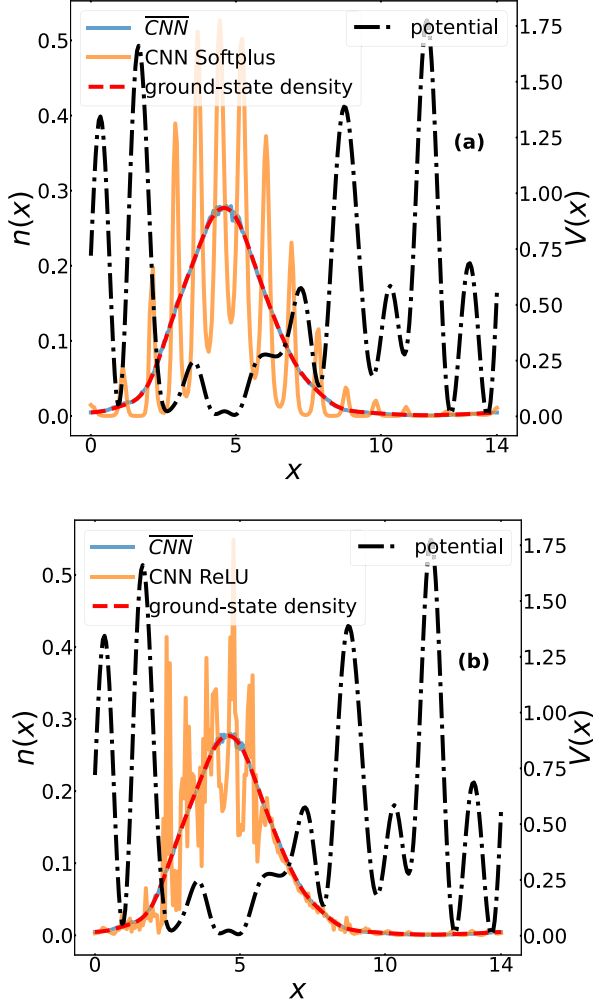


FIG. 1. Panels (a) and (b) show the external potentials $V(x)$ representing two instances of the optical speckle pattern (right vertical axis, the unit is the correlation energy E_c) and the corresponding ground-state density profiles (left vertical axis, units of $1/\gamma$). The spatial variable x is in units of the correlation length γ . In panel (a) the exact ground state is compared to the DFT result obtained using the standard CNN with Softplus activation function and using the $\overline{\text{CNN}}$. In panel (b) the standard CNN with ReLU activation is considered instead.

where N_{test} is the number of instances in the test set and $\sigma^2 = \frac{1}{N_{\text{test}}} \sum_{k=1}^{N_{\text{test}}} (t_{\text{gs},k} - \bar{t}_{\text{gs}})^2$ is the variance of their kinetic energies; with $\bar{t}_{\text{gs}} = \frac{1}{N_{\text{test}}} \sum_{k=1}^{N_{\text{test}}} t_{\text{gs},k}$ we denote the average kinetic energy. After training, the two variants of standard CNNs reach remarkable accuracies on the test set, meaning that, when they are provided with an exact ground-state density profile corresponding to a previously unseen speckle pattern, they accurately predict the associated ground-state kinetic energy and, via Eq. (3), also the total energy. The R^2 scores obtained with these two CNNs are reported in Table I. The R^2 scores reached by the $\overline{\text{CNN}}$ are comparable, but slightly inferior, to the ones obtained by the standard CNNs (see Table I). Remarkably, despite this (slightly) lower performance in kinetic energy predictions, the averaging operation drastically suppresses the noise in the functional derivative, allowing the

use of $\overline{\text{CNN}}$ in a gradient-based search of the ground-state energy and density profile. This is discussed in Sec. IV.

IV. RESULTS FOR GRADIENT-DESCENT OPTIMIZATION

To be suitable for the DFT framework, the deep-learned functional $\tilde{T}_\omega[n]$ should allow iterating the gradient descent process as long as is required to reach the minimum of $E[n] \equiv \tilde{T}_\omega[n] + \int dx V(x)n(x)$. Hereafter, we denote with $n_{\text{min}}(x)$ the density profile reached after gradient-descent optimization, and with $e_{\text{min}} = E[n_{\text{min}}(x)]$ the corresponding energy. The latter represents our estimate for the ground-state energy e_{gs} . Importantly, energies significantly lower than e_{gs} should never occur during the optimization process, as they would constitute a violation of the variational principle. As explained in Sec. II, the possible freezing in a local minimum significantly larger than e_{gs} can be circumvented by repeating the optimization from a different initial profile. Figure 1 displays the density profiles reached after $t_{\text{max}} = 10000$ steps of gradient descent for two representative instances of the speckle pattern. For these and all other results reported below, the learning rate used in gradient descent is $\eta = 10^{-3}$. Clearly, the standard CNNs lead to unphysical profiles, while the $\overline{\text{CNN}}$ provides an accurate approximation of the exact profile $n_{\text{gs}}(x)$. To shed light on this phenomenon, we analyze the behavior of the energy discrepancy

$$\Delta e = e_{\text{gs}} - e_{\text{min}} \quad (19)$$

and of the density discrepancy

$$|\Delta n| = \sqrt{\int dx [n_{\text{gs}}(x) - n_{\text{min}}(x)]^2}, \quad (20)$$

based on the L_2 metric

$$|n| = \sqrt{\int dx n^2(x)} \quad (21)$$

along the gradient-descent process. Specifically, we consider the average of the relative energy error $\langle \Delta e / e_{\text{gs}} \rangle$, of the relative absolute energy error $\langle |\Delta e| / e_{\text{gs}} \rangle$, and of the relative density error $\langle |\Delta n| / |n_{\text{gs}}| \rangle$ computed over a test set of 500 instances. Their dependence on the number of gradient-descent steps is shown in Fig. 2. The vertical bars indicate the standard deviation over this test set, meaning that they represent the fluctuations among different realizations of the speckle pattern. For both standard CNNs, after an initial decrease, the average absolute error increases. This means that the optimization process is not reliable, as it should be halted at an unknown intermediate number of steps. The average relative error becomes negative, indicating a violation of the variational principle. The density error also increases after many steps, corresponding to the formation of unphysical density profiles with large spurious spatial fluctuations, as exemplified in Fig. 1. Instead, the absolute relative energy error corresponding to the $\overline{\text{CNN}}$ (with $N_c = 260$ channels) systematically decreases until it saturates around a small value corresponding to $\sim 0.5\%$. The average error saturates close to $\langle \Delta e / e_{\text{gs}} \rangle \sim 0$, meaning that significant violation of the variational principle does not occur. The histograms shown




FIG. 2. Relative discrepancies of DFT predictions from exact ground-state results as a function of the number of steps t of the gradient-descent optimization. The results are averaged over a test set of 500 speckle-pattern instances. Three error metrics are shown: absolute energy discrepancy $\langle |\Delta e|/e \rangle$ [panel (a)], energy discrepancy $\langle \Delta e/e \rangle$ [panel (b)], and density discrepancy $\langle |\Delta n|/|n| \rangle$ [panel (c)]. The results of the standard CNN with ReLU activation function are compared to those of the $\overline{\text{CNN}}$ ($N_c = 260$). The vertical bars represent the standard deviation over the test set.

in Fig. 3 compare the energy and the density errors obtained with the $\overline{\text{CNN}}$ and with the standard CNN with ReLU activation function (this outperforms the corresponding model with Softplus activation) after $t_{\max} = 10\,000$ steps of the gradient descent optimization. The $\overline{\text{CNN}}$ energies are concentrated around zero error, while the standard CNN results are broadly distributed in the region of negative energy errors, corresponding to strong violations of the variational principle.

Notably, the energy predictions obtained by performing gradient-descent optimization with the $\overline{\text{CNN}}$ systematically improve when the number of channels N_c increases. This effect is shown in Fig. 4. Notice that, for small N_c , small violations of the variational principle still occur in rare cases. However, they vanish for larger N_c . In fact, the average absolute energy discrepancy obtained after gradient descent with the $N_c = 260$ $\overline{\text{CNN}}$ is $\langle |\Delta e|/e \rangle \simeq 0.2\%$, which is approximately twice the corresponding discrepancy obtained on a test set of exact ground-state profiles $n_{\text{gs}}(x)$, namely

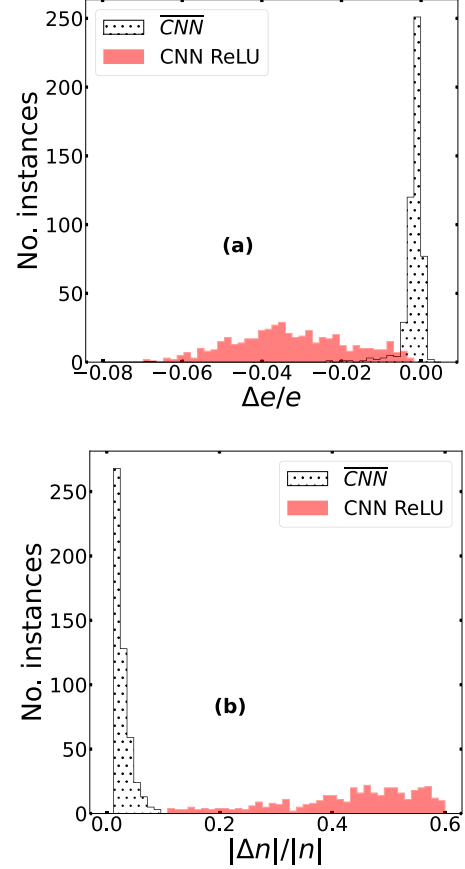


FIG. 3. Histograms of the relative energy discrepancies $\Delta e/e$ [panel (a)] and the density discrepancies $|\Delta n|/|n|$ [panel (b)] for a test set of 500 speckle-pattern instances, after $t_{\max} = 10\,000$ steps of the gradient-descent optimization. The results of the standard CNN with ReLU activation function are compared to those of the $\overline{\text{CNN}}$.

$\langle |\Delta e_{\text{ML}}|/e \rangle \simeq 0.1\%$. Notice that this approximate doubling effect is expected, since the former error is also affected by the approximation in the density profile, while the latter corresponds to the $\overline{\text{CNN}}$ prediction on the exact profile. This means that gradient-descent optimization successfully identifies the ground state, within the residual uncertainty of the ML model. Increasing N_c leads also to more accurate density profiles [see panel (b) of Fig. 4] and to the reduction of the spatial noise observed in the results provided by standard CNN (see Fig. 1 and also Refs. [10,18]).

To quantify this spatial noise, we consider the following metric:

$$\Delta A = \frac{1}{N_g} \left\langle \sum_{i=1}^{N_g} \left(\left| \frac{\nabla n_{\min}(x_i)}{\nabla n_{\text{gs}}(x_i)} \right| - 1 \right) \right\rangle. \quad (22)$$

It measures the error in the derivative of the density profile. Inaccurate profiles are characterized by large positive values $\Delta A \gg 1$, due to spurious spatial fluctuations, while exact predictions lead to $\Delta A = 0$. We find that large N_c values lead to accurate and smooth density profiles [see panel (c) of Fig. 4], indicating the effectiveness of the average-channel layer. Residual local spurious fluctuations in the density profiles might be further suppressed via filtering procedures in




FIG. 4. Histograms of the relative energy discrepancies $\Delta e/e$ [panel (a)], of the density discrepancies $|\Delta n|/|n|$ [panel (b)], and of the noise metric ΔA defined in Eq. (22) [panel (c)], for a test set of 500 speckle-pattern instances, after $t_{\max} = 10\,000$ steps of the gradient descent optimization. The results of the $\overline{\text{CNN}}$ s with different number of convolutional channels N_c are shown.

postprocessing. Still, accurate DFT predictions are essential to avoid introducing biases by strong filtering procedures.

It is worth further emphasizing that using the standard metrics of deep learning on training and test sets of exact density profiles is not necessarily helpful to predict the performance of the ML functional in the variational minimization of DFT. This is exemplified by the scatter plots of Fig. 5. They display the errors obtained after gradient descent optimization versus




FIG. 5. Absolute relative discrepancy $\langle |\Delta e|/|e| \rangle$ [panel (a)] and scatter plot of density discrepancy $\langle |\Delta n|/|n| \rangle$ [panel (b)] after gradient-descent optimization ($t_{\max} = 10\,000$, average over 500 instances) vs coefficient of determination R^2 over a test set of exact ground-state densities. Three neural networks are compared: standard CNN with ReLU and with Softplus activation functions, and $\overline{\text{CNN}}$.

the coefficient of determination Eq. (18) computed on a test set of exact ground-state density profiles. The two standard CNNs and the $\overline{\text{CNN}}$ with three values of N_c are considered. No (anti)correlation is clearly noticeable, meaning that high prediction accuracies on exact ground states do not necessarily correspond to highly effective functionals for DFT. This should be taken into account in the future development of deep learning techniques for DFT.

V. CONCLUSIONS

The progress in databased DFT is currently being hindered by the instabilities encountered when using ML functionals in gradient-based optimization. We presented a promising approach to circumvent this problem. This relies on the implementation of a deep neural network tailored to DFT. Specifically, we have shown that when an interchannel averaging layer is included, beyond the standard convolutional, pooling, and dense layers, gradient-descent optimization can be iterated at will, obtaining accurate ground-state energies and density profiles and avoiding violations of the variational principle beyond residual uncertainties from the imperfect training of the regression model. Our analysis has focused on a realistic one-dimensional model for noninteracting atoms in optical speckle disorder, which leads to rather variegated density profiles compared to models addressed in previous

studies on ML-based DFT. For completeness, in Appendix the performance of the CNN in Gaussian-well models borrowed from the literature is demonstrated. On the one hand, this further analysis indicates the rather general range of applicability of our tailored neural network. On the other hand, it points out the need of training sets including significantly variegated density profiles. To favor future comparative studies, our training datasets are made freely available at Ref. [36].

Additional challenges are going to be faced in the further development of ML techniques for DFT [37]. A further assessment of the network effectiveness should focus on higher dimensional and interacting models. In the DFT formalism, moving from single-particle to many-body problems is less challenging than in wave-function-based methods, since observables are still obtained from the single-particle density. Therefore, we expect the CNN to be useful also in the many-body context. Clearly, generating training datasets is, in that case, more demanding [19,38]. The learning process has to be accelerated, and the following strategies could be adopted. Incorporating physics knowledge into the deep-learning framework is a possible strategy [39,40]. Another promising approach is transfer learning. This technique has already proven suitable to accelerate the supervised learning of the ground-state properties of both noninteracting and interacting quantum systems [41–43]. Interestingly, even extrapolations were proven feasible, meaning that (scalable) networks trained on relatively small systems provided accurate predictions for larger sizes or larger particle numbers. These techniques might be adopted also in the framework of DFT. We leave these endeavors to future investigations.

ACKNOWLEDGMENTS

This work was supported by the Italian Ministry of University and Research under the PRIN2017 project CEnTraL 20172H2SC4. S.P. acknowledges PRACE for awarding access to the Fenix Infrastructure resources at Cineca, which are partially funded by the European Union’s Horizon 2020 research and innovation program through the ICEI project under Grant Agreement No. 800858. S.P. also acknowledges the Cineca award under the ISCRA initiative for the availability of high-performance computing resources and support.

APPENDIX: GAUSSIAN-WELL POTENTIALS

To favor direct comparison, and to further characterize the effectiveness of the CNN, here we consider a different testbed model, borrowed from previous studies. It describes narrow potential wells that confine the particle in the central region of a finite box. Setting the units such that $\hbar^2/(m\ell^2) = 1$, where ℓ is the length unit, the (adimensional) single-particle Hamiltonian reads

$$H = -\frac{1}{2} \frac{d^2}{dx^2} + V(x), \quad (\text{A1})$$

with the external potential

$$V(x) = -\sum_{i=1}^3 a_i \exp\left[-\frac{(x-b_i)^2}{2c_i^2}\right]. \quad (\text{A2})$$

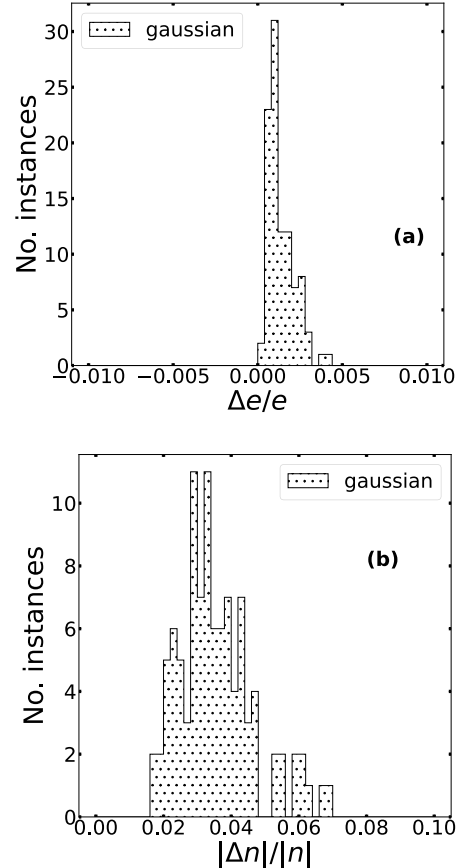


FIG. 6. Histograms of the relative energy discrepancies $\Delta e/e$ [panel (a)] and the density discrepancies $|\Delta n|/|n|$ [panel (b)] for a test set of 100 Gaussian-well potentials, after $t_{\max} = 10\,000$ steps of the gradient-descent optimization.

The model parameters a_i , b_i , and c_i are sampled from uniform probability distributions in the following ranges: $a_i \in [1, 40]$, $b_i \in [1.8, 4.2]$, and $c_i \in [0.12, 0.4]$. The box size is $L = 6$, and hard-wall boundary conditions are adopted. The chosen box is $L = 6$, rather than $L = 1$ as in Ref. [18], so that the ground-state density profiles essentially vanish before reaching the boundaries. This strongly suppresses the role of the choice of the boundary conditions. Indeed, we find negligible variations in the ground-state energies with hard-wall compared to periodic boundary conditions. With the size $L = 1$, the boundary effects are sizable, and the density profiles corresponding to different potential instances display small variations when hard-wall boundaries are adopted. This does not allow an effective training of the deep neural network, reintroducing instabilities in the gradient-descent optimization. We find that this problem is solved either enlarging the system size, e.g., to $L = 6$, as shown hereafter, or by adopting periodic boundary conditions in a small box, as mentioned later on in this paragraph.

A CNN is trained on a global dataset of 150 000 instances (90% train, 10% validation), using the same structure (2 blocks, pooling size [4,2], kernel size 13, and 260 hidden channels) and the same hyperparameters as in the case of the speckle potential with a number of epochs $N_e = 3000$. However, here the convolutional layers use zero padding,

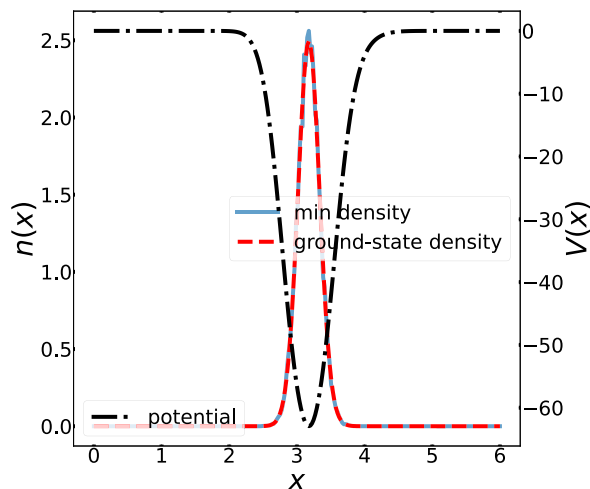


FIG. 7. External potential $V(x)$ corresponding to an instance of the Gaussian-well potential (right vertical axis, unit of $\frac{\hbar^2}{m\ell^2}$) and the corresponding ground-state density profiles (left vertical axis, units of $1/\ell$). ℓ is the length unit, used on the horizontal axis. The density configuration obtained by gradient-descent minimization is compared to the exact density profile.

as opposed to the periodic padding adopted for the speckle potentials, which are defined within periodic boundary conditions. The gradient descent performance on a test set of 100 instances is visualized in Fig. 6. Again, we find that gradient-descent optimization ($t_{\max} = 10\,000$) leads to accurate ground-state energies ($\langle |\Delta e|/e \rangle \simeq 0.14\%$) and density profiles ($\langle |\Delta n|/n \rangle \simeq 3\%$), without sizable violation of the variational condition. For illustrative purposes, an example of Gaussian-well potential, with the corresponding density

profile obtained with $L = 6$, is shown in Fig. 7. To facilitate comparison, we report the model details and the performance metrics using the units and the conventions adopted in some previous studies. These considered one-electron models in atomic units, so that ℓ corresponds to the Bohr radius, and the energy unit $\hbar^2/(m\ell^2)$ corresponds to one Hartree (Ha). The mean kinetic energy of the $L = 6$ model is $\bar{t}_{\text{gs}} \simeq 3.994$ Ha and the standard deviation is quite large, namely $\sigma \simeq 1.457$ Ha, with maximum value 12.957 Ha and minimum value 0.2605 Ha, corresponding to considerably variegated ground states. The average kinetic-energy prediction error on a test set of $N_{\text{test}} = 15\,000$ instances is $\langle |\Delta t_{\text{ML}}| \rangle \equiv N_{\text{test}}^{-1} \sum_{k=1}^{N_{\text{test}}} |t_{\text{gs},k} - \tilde{T}_{\omega}[n_{\text{gs},k}]| \simeq 8 \times 10^{-4}$ Ha; the standard deviation of $|\Delta t_{\text{ML}}|$ is $\sigma \simeq 11 \times 10^{-4}$ Ha. After gradient descent, the average absolute kinetic-energy error is $\langle |\tilde{T}_{\omega}[n_{\text{min}}] - t_{\text{gs}}| \rangle \simeq 0.0171$ Ha and the corresponding standard deviation is $\sigma \simeq 0.0191$ Ha. The average absolute density discrepancy is $\langle |\Delta n| \rangle = 0.031$ and the corresponding standard deviation is $\sigma \simeq 0.008$.

It is also worth mentioning that, considering $L = 1$ and periodic boundary conditions, the accuracy metrics are comparable to those mentioned above, namely $\langle |\Delta e|/e \rangle \simeq 0.3\%$ and $\langle |\Delta n|/n \rangle \simeq 3\%$. In this case, periodic padding is used, and the first left and right periodic images of the Gaussian wells are included to make the potential $V(r)$ essentially periodic. The potential parameters are sampled in the following ranges: $b_i \in [0.1, 0.9]$, $a_i \in [1, 40]$, and $c_i \in [0.03, 0.1]$. When combined with periodic boundary conditions, these allow us to create sufficiently variegated density profiles despite the small system size. These findings indicate that the CNN represents a flexible regression model to apply ML-based DFT to rather arbitrary external potentials, whereby the density profiles display significant variations among different random realizations of the sample.

- [1] K. Burke, Perspective on density functional theory, *J. Chem. Phys.* **136**, 150901 (2012).
- [2] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
- [3] M. Levy, Universal variational functionals of electron densities, first-order density matrices, and natural spin-orbitals and solution of the v -representability problem, *Proc. Natl. Acad. Sci. (USA)* **76**, 6062 (1979).
- [4] E. H. Lieb, Density functionals for Coulomb systems, in *Inequalities: Selecta of Elliott H. Lieb*, edited by M. Loss and M. B. Ruskai (Springer, Berlin, 2002), pp. 269–303
- [5] A. J. Cohen, P. Mori-Sánchez, and W. Yang, Challenges for density functional theory, *Chem. Rev.* **112**, 289 (2012).
- [6] V. Dunjko and H. J. Briegel, Machine learning & artificial intelligence in the quantum domain: A review of recent progress, *Rep. Prog. Phys.* **81**, 074001 (2018).
- [7] 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).
- [8] J. Carrasquilla, Machine learning for quantum matter, *Adv. Phys.-X* **5**, 1797528 (2020).
- [9] H. Kulik, T. Hammerschmidt, J. Schmidt, S. Botti, M. A. L. Marques, M. Boley, M. Scheffler, M. Todorović, P. Rinke, C. Oses, A. Smolyanyuk, S. Curtarolo, A. Tkatchenko, A. Bartok, S. Manzhos, M. Ihara, T. Carrington, J. Behler, O. Isayev, M. Veit *et al.*, Roadmap on machine learning in electronic structure, *Electron. Struct.* (to be published).
- [10] J. C. Snyder, M. Rupp, K. Hansen, K.-R. Müller, and K. Burke, Finding Density Functionals with Machine Learning, *Phys. Rev. Lett.* **108**, 253002 (2012).
- [11] L. Li, J. C. Snyder, I. M. Pelaschier, J. Huang, U.-N. Niranjan, P. Duncan, M. Rupp, K.-R. Müller, and K. Burke, Understanding machine-learned density functionals, *Int. J. Quantum Chem.* **116**, 819 (2016).
- [12] F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, Bypassing the Kohn-Sham equations with machine learning, *Nat. Commun.* **8**, 872 (2017).
- [13] K. Ryczko, D. A. Strubbe, and I. Tamblin, Deep learning and density-functional theory, *Phys. Rev. A* **100**, 022512 (2019).
- [14] J. Kirkpatrick, B. McMorrow, D. H. P. Turban, A. L. Gaunt, J. S. Spencer, A. G. D. G. Matthews, A. Obika, L. Thiry, M. Fortunato, D. Pfau, L. R. Castellanos, S. Petersen,

- A. W. R. Nelson, P. Kohli, P. Mori-Sánchez, D. Hassabis, and A. J. Cohen, Pushing the frontiers of density functionals by solving the fractional electron problem, *Science* **374**, 1385 (2021).
- [15] J. Nelson, R. Tiwari, and S. Sanvito, Machine learning density functional theory for the Hubbard model, *Phys. Rev. B* **99**, 075132 (2019).
- [16] J. R. Moreno, G. Carleo, and A. Georges, Deep Learning the Hohenberg-Kohn Maps of Density Functional Theory, *Phys. Rev. Lett.* **125**, 076402 (2020).
- [17] M. M. Denner, M. H. Fischer, and T. Neupert, Efficient learning of a one-dimensional density functional theory, *Phys. Rev. Res.* **2**, 033388 (2020).
- [18] R. Meyer, M. Weichselbaum, and A. W. Hauser, Machine learning approaches toward orbital-free density functional theory: Simultaneous training on the kinetic energy density functional and its functional derivative, *J. Chem. Theory Comput.* **16**, 5685 (2020).
- [19] K. Ryczko, S. J. Wetzel, R. G. Melko, and I. Tamblin, Toward orbital-free density functional theory with small data sets and deep learning, *J. Chem. Theory Comput.* **18**, 1122 (2022).
- [20] J. C. Snyder, M. Rupp, K.-R. Müller, and K. Burke, Nonlinear gradient denoising: Finding accurate extrema from inaccurate functional derivatives, *Int. J. Quantum Chem.* **115**, 1102 (2015).
- [21] L. Li, T. E. Baker, S. R. White, and K. Burke, Pure density functional for strong correlation and the thermodynamic limit from machine learning, *Phys. Rev. B* **94**, 245129 (2016).
- [22] In interacting systems, the unknown functional would include also correlation effects, while the mean-field contribution would be conveniently encoded in the so-called Hartree term.
- [23] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, Automatic differentiation in PyTorch, in *Proceedings of the NIPS 2017 Autodiff Workshop The Future of Gradient-based Machine Learning Software and Techniques* (Long Beach, CA, US, 2017).
- [24] W. C. Witt, B. G. del Rio, J. M. Dieterich, and E. A. Carter, Orbital-free density functional theory for materials research, *J. Mater. Res.* **33**, 777 (2018).
- [25] J. Billy, V. Josse, Z. Zuo, A. Bernard, B. Hambrecht, P. Lugan, D. Clément, L. Sanchez-Palencia, P. Bouyer, and A. Aspect, Direct observation of Anderson localization of matter waves in a controlled disorder, *Nature (London)* **453**, 891 (2008).
- [26] G. Roati, C. D'Errico, L. Fallani, M. Fattori, C. Fort, M. Zaccanti, G. Modugno, M. Modugno, and M. Inguscio, Anderson localization of a non-interacting Bose-Einstein condensate, *Nature (London)* **453**, 895 (2008).
- [27] J. W. Goodman, *Speckle Phenomena in Optics: Theory and Applications* (Roberts and Company, Greenwood Village, CO, United States, 2007).
- [28] T. Prat, N. Cherroret, and D. Delande, Semiclassical spectral function and density of states in speckle potentials, *Phys. Rev. A* **94**, 022114 (2016).
- [29] S. Pilati and P. Pieri, Supervised machine learning of ultracold atoms with speckle disorder, *Sci. Rep.* **9**, 5613 (2019).
- [30] P. Mujal, A. Polls, S. Pilati, and B. Juliá-Díaz, Few-boson localization in a continuum with speckle disorder, *Phys. Rev. A* **100**, 013603 (2019).
- [31] E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan, Scaling Theory of Localization: Absence of Quantum Diffusion in Two Dimensions, *Phys. Rev. Lett.* **42**, 673 (1979).
- [32] P. W. Anderson, Absence of diffusion in certain random lattices, *Phys. Rev.* **109**, 1492 (1958).
- [33] J. Huntley, Speckle photography fringe analysis: assessment of current algorithms, *Appl. Opt.* **28**, 4316 (1989).
- [34] M. Modugno, Collective dynamics and expansion of a Bose-Einstein condensate in a random potential, *Phys. Rev. A* **73**, 013606 (2006).
- [35] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [36] E. Costa, G. Scriva, S. Pilati, and R. Fazio, Data for deep learning density functionals for gradient descent optimization, Zenodo (2022); doi: [10.5281/zenodo.7108256](https://doi.org/10.5281/zenodo.7108256).
- [37] R. Pederson, B. Kalita, and K. Burke, Machine learning and density functional theory, *Nat. Rev. Phys.* **4**, 357 (2022).
- [38] K. Ryczko, J. T. Krogel, and I. Tamblin, Machine learning diffusion monte carlo energy densities, [arXiv:2205.04547](https://arxiv.org/abs/2205.04547).
- [39] R. Nagai, R. Akashi, and O. Sugino, Completing density functional theory by machine learning hidden messages from molecules, *npj Comput. Mater.* **6**, 43 (2020).
- [40] L. Li, S. Hoyer, R. Pederson, R. Sun, E. D. Cubuk, P. Riley, and K. Burke, Kohn-Sham Equations as Regularizer: Building Prior Knowledge into Machine-Learned Physics, *Phys. Rev. Lett.* **126**, 036401 (2021).
- [41] K. Mills, K. Ryczko, I. Luchak, A. Domurad, C. Beeler, and I. Tamblin, Extensive deep neural networks for transferring small scale learning to large scale systems, *Chem. Sci.* **10**, 4129 (2019).
- [42] N. Saraceni, S. Cantori, and S. Pilati, Scalable neural networks for the efficient learning of disordered quantum systems, *Phys. Rev. E* **102**, 033301 (2020).
- [43] P. Mujal, A. Martínez Miguel, A. Polls, B. Juliá-Díaz, and S. Pilati, Supervised learning of few dirty bosons with variable particle number, *SciPost Phys.* **10**, 073 (2021).