



Neural functional theory for inhomogeneous fluids: Fundamentals and applications

Florian Sammüller^{a,1} , Sophie Hermann^a , Daniel de las Heras^a , and Matthias Schmidt^{a,1}

Edited by Robert Evans, University of Bristol, Bristol, United Kingdom; received July 21, 2023; accepted October 7, 2023 by Editorial Board Member Daan Frenkel

We present a hybrid scheme based on classical density functional theory and machine learning for determining the equilibrium structure and thermodynamics of inhomogeneous fluids. The exact functional map from the density profile to the one-body direct correlation function is represented locally by a deep neural network. We substantiate the general framework for the hard sphere fluid and use grand canonical Monte Carlo simulation data of systems in randomized external environments during training and as reference. Functional calculus is implemented on the basis of the neural network to access higher-order correlation functions via automatic differentiation and the free energy via functional line integration. Thermal Noether sum rules are validated explicitly. We demonstrate the use of the neural functional in the self-consistent calculation of density profiles. The results outperform those from state-of-the-art fundamental measure density functional theory. The low cost of solving an associated Euler–Lagrange equation allows to bridge the gap from the system size of the original training data to macroscopic predictions upon maintaining near-simulation microscopic precision. These results establish the machine learning of functionals as an effective tool in the multiscale description of soft matter.

classical density functional theory | machine learning | statistical mechanics | soft matter

The problem with density functional theory (DFT) is that you do not know the density functional. Although this quip by the late and great Yasha Rosenfeld (1) was certainly meant in jest to a certain degree, it does epitomize a structural assessment of classical DFT (2–5). As a general formulation of many-body statistical physics, the framework comprises a beautiful and far-reaching skeleton of mathematical formalism centered around a formally exact variational minimization principle (2, 6). In practice, however, the theory needs to be fleshed out by approximations of all means conceivable in our efforts to get to grips with the coupled many-body problem that is under consideration. Specifically, it is the excess (over ideal gas) intrinsic Helmholtz free energy $F_{\text{exc}}[\rho]$, expressed as a functional of the position-resolved density profile $\rho(\mathbf{r})$, which needs to be approximated.

Decades of significant theoretical efforts have provided us with a single exact functional, that for nonoverlapping hard rods in one spatial dimension, as obtained by another hero in the field, Jerry Percus (7). Nevertheless, useful DFT approximations range from the local density approximation for large-scale features which are decoupled from microscopic length scales, to square-gradient functionals with their roots in the 19th century, to the arguably most important modern development, that of the fundamental measure theory (FMT) as kicked off by Rosenfeld (8) and much refined ever since (9–16). FMT is a geometry-based framework for the description of hard sphere systems and it has deep roots in the Percus–Yevick (17) and scaled-particle theories (4), which Rosenfeld was able to unify and generalize based on his unique theoretical insights (18).

The realm of soft matter (19–21) stretches far beyond the hard sphere fluid. FMT remains relevant though in the description of a reference system as used, e.g., in studies of hydrophobicity, where the behavior of realistic water models (22, 23) is traced back to the simpler Lennard-Jones fluid, which in turn is approximated via the hard sphere FMT functional plus a mean-field contribution for describing interparticle attraction (20, 24, 25). Further topical uses of FMT include the analysis of the three-dimensional electrolyte structure near a solid surface (26, 27) and the problem of the decay length of correlations in electrolytes (28).

There is a current surge in the use of machine learning techniques in soft matter, e.g., for its characterization (29), engineering of self-assembly (30), structure detection (31), and for learning many-body potentials (32, 33). Within classical DFT, machine learning was used to address ordering of confined liquid crystals (34), and free energy functionals were

Significance

Classical density functional theory (DFT) is a powerful approach for predicting inhomogeneous fluid equilibria. The development of functionals that model the rich interparticle correlations remains a major theoretical challenge though. Here, we introduce a universal machine learning scheme for training a neural network with Monte Carlo simulation data in order to represent the density functional relationship via direct one-body correlations. We show for the well-studied hard sphere fluid that the resulting neural correlation functional exceeds DFT in accuracy all current analytic approaches and that it provides immediate access to the full structure of homogeneous and heterogeneous fluid states. The corresponding neural DFT yields density profiles at near-simulation precision, and it applies directly to multiscale problems at low computational cost.

Author contributions: F.S., D.d.l.H., and M.S. designed research; F.S. and S.H. performed research; F.S. analyzed data; and F.S., S.H., D.d.l.H., and M.S. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. R.E. is a guest editor invited by the Editorial Board.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: florian.sammuller@uni-bayreuth.de or Matthias.Schmidt@uni-bayreuth.de.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2312484120/-/DCSupplemental>.

Published December 7, 2023.

obtained for one-dimensional systems from convolutional (35) and equation-learning (36) networks as well as within a Bayesian inference approach (37). Cats et al. (38) used machine learning to improve the standard mean-field approximation of the excess Helmholtz free-energy functional for the Lennard-Jones fluid. In nonequilibrium, de las Heras et al. (39) have reported a method to machine-learn the functional relationship of the local internal force for a steady uniaxial compressional flow of a Lennard-Jones fluid at constant temperature. As prescribed by power functional theory (40, 41), the functional dependence in nonequilibrium not only incorporates the density profile but also the one-body current.

In this work, we return to the problem of describing and predicting the structure and thermodynamics of inhomogeneous equilibrium fluids. We show that a neural network can be trained to accurately represent the functional dependence of the one-body direct correlation function with respect to the density profile. The presented methods are directly applicable to virtually arbitrary fluids with short-ranged interparticle interactions. In the following, we focus on the well-studied hard sphere fluid in order to exemplify our framework and to challenge the available highly accurate analytic approaches from liquid integral equation theory and FMT. We give more details about the feasibility of generalizations in the discussion. Reference data for training and testing the model are provided by grand canonical Monte Carlo (GCMC) simulations that cover a broad range of randomized inhomogeneous environments in planar geometry.

We implement functional calculus on the basis of the trained neural functional to infer related physical quantities and demonstrate their consistency with known literature results both in bulk and in inhomogeneous systems. In particular, we highlight the accessibility of the fluid pair structure, the determination of free energies and equations of state as well as the validation of thermal Noether sum rules (42). These results corroborate that the neural functional exceeds its role as a mere interpolation device and instead possesses significant representational power as a genuine density functional for the prediction of nontrivially related physical properties. We apply the trained neural network in the DFT Euler–Lagrange equation, which enables the self-consistent calculation of density profiles and which hence constitutes a neural-network-based DFT or short neural DFT. This method alleviates conventional DFT from the burden of having to find suitable analytic approximations while still surpassing even the most profound existing treatments of the considered hard sphere fluid via FMT functionals (8, 13, 14) in accuracy. We further demonstrate the fitness of the method for the straightforward application to multiscale problems. Neural DFT therefore provides a way to transfer near-simulation microscopic precision to macroscopic length scales, which serves as a technique to predict properties of inhomogeneous systems which far exceed typical box sizes of the original training data.

This work is structured as follows. The relevant physical background of liquid state theory is provided in Section A.1. Details of the simulations as well as of the neural network are given in Sections A.2 and A.3. The training procedure and results for the achieved metrics that measure its convergence are presented in Section A.4. We proceed by testing physical properties of the trained model and use automatic differentiation of the neural network in Section B.1 to access pair correlations, which are then compared to bulk results from both the Percus–Yevick theory and from simulations. The consistency of the neural direct correlation functional to satisfy thermal Noether sum rules is validated in Section B.2, and different ways to obtain the bulk equation

of state as well as free energies in inhomogeneous systems are given in Section B.3. In Section C.1, we show the application of the neural functional to the self-consistent calculation of density profiles via the DFT Euler–Lagrange equation and describe the technical details and conceptual advantages of this neural DFT over analytic approaches. In Section C.2, the results are compared to those from FMT, and in Section C.3, the relevance of the method for making macroscopic predictions is illustrated for cases of randomized external potential and for sedimentation between hard walls on length scales that far exceed the training simulation box sizes. We conclude with a discussion of the results and give an outlook on possible improvements and extensions of the method as well as to its application for different fluid types, in more general geometries and in nonequilibrium.

Results

A. Machine Learning Intrinsic Correlations.

A.1. Physical background. We start with the standard relation for the one-body direct correlation function $c_1(\mathbf{r})$ of liquid state theory (4),

$$c_1(\mathbf{r}) = \ln \rho(\mathbf{r}) + \beta V_{\text{ext}}(\mathbf{r}) - \beta \mu, \quad [1]$$

where \mathbf{r} denotes the spatial position and $\beta = 1/(k_B T)$ with the Boltzmann constant k_B and absolute temperature T . The three terms on the right-hand side of Eq. 1 represent respectively the ideal gas contribution, the external potential $V_{\text{ext}}(\mathbf{r})$ and the influence of the particle bath at chemical potential μ . The logarithm in Eq. 1 is understood as $\ln[\Lambda^3 \rho(\mathbf{r})]$ with the thermal wavelength Λ , which can be set to the particle size σ without any loss of information in the present classical context. For a prescribed external potential $V_{\text{ext}}(\mathbf{r})$, knowledge of the corresponding equilibrium density profile $\rho(\mathbf{r})$ allows to compute $c_1(\mathbf{r})$ explicitly via Eq. 1. This relationship can be viewed as a locally resolved chemical potential balance: the contribution from the ideal gas, $k_B T \ln \rho(\mathbf{r})$, from the external potential, $V_{\text{ext}}(\mathbf{r})$, and from interparticle interactions, $-k_B T c_1(\mathbf{r})$, add up at each position to μ , which is necessarily uniform throughout an equilibrium system.

However, the notation in Eq. 1 is oblivious to a central result shown by Evans (2), thereby kicking off a modern theory for the description of inhomogeneous fluids. For given type of internal interactions, the spatial variation of the function $c_1(\mathbf{r})$ is already uniquely determined by the spatial form of the density profile $\rho(\mathbf{r})$ alone, without the need to invoke the external potential explicitly. From this vantage point of classical DFT, the dependence of $c_1(\mathbf{r})$ on $\rho(\mathbf{r})$ is not merely pointwise but rather with respect to the values of the entire density profile, which determine $c_1(\mathbf{r})$ at each given position \mathbf{r} . Formally, this relationship is exact (2, 4), and it constitutes a functional dependence $c_1(\mathbf{r}; [\rho])$, which is indicated by brackets here and in the following and which is in general nonlinear and nonlocal. As we will demonstrate, the existence of such a universal functional mapping makes the problem of investigating inhomogeneous fluids particularly amenable to supervised machine learning techniques.

In most formulations of classical DFT, one exploits the fact that the intrinsic excess free energy functional $F_{\text{exc}}[\rho]$ acts as a functional generator such that the one-body direct correlation function is obtained via functional differentiation with respect to the density profile,

$$c_1(\mathbf{r}; [\rho]) = -\frac{\delta \beta F_{\text{exc}}[\rho]}{\delta \rho(\mathbf{r})}. \quad [2]$$

A compact description of standard formulae for the calculation of functional derivatives can be found in ref. 41. In order to make progress in concrete applications, one typically needs to rely on using an approximate form of $F_{\text{exc}}[\rho]$ for the specific model under consideration, as determined by its interparticle interactions. DFT is a powerful framework, as using $c_1(\mathbf{r}; [\rho])$ obtained from Eq. 2 with a suitable expression for $F_{\text{exc}}[\rho]$ turns Eq. 1 into an implicit equation for the equilibrium density profile $\rho(\mathbf{r})$. In the presence of a known form of $V_{\text{ext}}(\mathbf{r})$, one can typically solve Eq. 1 very efficiently, allowing ease of parameter sweeps, e.g., for exhaustive phase diagram explorations. On the downside, $F_{\text{exc}}[\rho]$ and thus also $c_1(\mathbf{r}; [\rho])$ remain approximate, and the development of analytic tools has certainly slowed down over several years if not decades.

Here, we proceed differently and bypass the excess free energy functional $F_{\text{exc}}[\rho]$ at first. Instead, we use a deep neural network to learn and to represent the functional relationship $\rho(\mathbf{r}) \rightarrow c_1(\mathbf{r})$ directly, which has significant advantages both for the generation of suitable training data as well as for the applicability of the model in the determination of fluid equilibria. This investigation is based on GCMC simulations that serve to provide training, validation, and test data. Discriminating between these three roles of use is standard practice in machine learning and we give further details below.

A.2. Simulation method. Generating the simulation data is straightforward, and we use the following strategy, adopted to planar situations where the position dependence is on a single position variable x while the system remains translationally invariant in the y and z direction. This geometry is highly relevant to identify the physics in planar capillary and adsorption situations and facilitates ease of accurate sampling. We employ randomized simulation conditions by generating external potentials of the form

$$V_{\text{ext}}(x) = \sum_{n=1}^4 A_n \sin\left(\frac{2\pi nx}{L} + \phi_n\right) + \sum_n V_n^{\text{lin}}(x), \quad [3]$$

where A_n and ϕ_n are randomly selected Fourier coefficients and phases, respectively, and L is the simulation box length in the x direction. The phases ϕ_n are chosen uniformly in the interval $[0, 2\pi)$, and values of A_n are drawn from a normal distribution with zero mean and variance 2.5. We choose $L = 20\sigma$, although there is no specific compliance requirement for the neural network (see below), and the lateral box lengths are set to 10σ to minimize finite-size effects. Periodic boundary conditions apply in all spatial directions. The sinusoidal terms in $V_{\text{ext}}(x)$ are complemented by up to five piecewise linear functions $V_n^{\text{lin}}(x) = V_1 + (V_2 - V_1)(x - x_1)/(x_2 - x_1)$ for $x_1 < x < x_2$ and 0 otherwise, for which the parameters $0 < x_1 < x_2 < L$, V_1 , and V_2 are again chosen randomly. The locations x_1 and x_2 are distributed uniformly while V_1 and V_2 follow again from an unbiased normal distribution with variance 4. Additionally to the discontinuous linear segments, we explicitly impose planar hard walls in a subset of the simulations by setting $V_{\text{ext}}(x) = \infty$ for $x < x_w/2$ and $x > L - x_w/2$, i.e., near the borders of the simulation domain; the width x_w of the wall is chosen randomly in the interval $1 \leq x_w/\sigma \leq 3$. To cover a broad range from dilute to dense systems, the chemical potential is chosen uniformly within the range $-5 \leq \beta\mu \leq 10$ for each respective GCMC simulation run. The observed mean densities range from $0.006\sigma^{-3}$ to $0.803\sigma^{-3}$, yet smaller and much larger local densities occur due to the inhomogeneous nature of the systems.

In total, 750 such GCMC runs are used, where for given form of $V_{\text{ext}}(x)$, the planar one-body profiles $\rho(x)$ and $c_1(x)$ are

obtained. The former is acquired from straightforward histogram filling and the latter from evaluating Eq. 1 on the basis of the sampled histogram for $\rho(x)$ as well as the known form of $V_{\text{ext}}(x)$ and value of μ for the specific run under consideration. As Eq. 1 is undefined for vanishing density, we have excluded regions where $\rho(x) = 0$ such as within the hard walls. By modern standards of computational resources, the workload for the generation of the simulation data is only moderate at a total CPU time of $\sim 10^4$ h.

A.3. Neural network. We use a deep neural network (43) to represent the functional map from the density profile to the local value of the one-body direct correlation function at a given point. That is, instead of the entire function, we construct the network to output only the scalar value $c_1(x)$ for a certain position x when supplied with the surrounding inhomogeneous density. The relevant section of the density profile comprises the values of $\rho(x)$ in a specified window around a considered location x , as described below. Despite the locality of the method, access to the entire (discretized) one-body direct correlation profile is immediate via evaluation of the neural network at pertinent positions x across the domain of interest. Multiple local evaluations of the network remain performant on highly parallel hardware such as GPUs when passing the input accordingly in batches. A schematic picture of the network architecture is given in Fig. 1 and is explained in the following.

The functional dependence on the density profile is realized by providing discretized values of $\rho(x)$ on an equidistant grid with resolution $\Delta x = 0.01\sigma$. As $c_1(x; [\rho])$ depends only on the immediately surrounding density profile around a fixed location x , we restrict the input range x' to a sufficiently large window $x' \leq |x - x_c|$. We choose the cutoff $x_c = 2.56\sigma$ based on simulation data for the bulk direct correlation function (44) and on the evaluation of training metrics for different window sizes x_c . Increasing the value of x_c further led to no improvement in the performance of the trained neural network.

This behavior is expected from theoretical considerations, as the one-body direct correlation function vanishes quickly for short-ranged pair potentials (4). We recall that in FMT, $x_c = \sigma$ by construction. Note that the choice of $c_1(x; [\rho])$ as our target functional is not coincidental but that its quick spatial decay rather is a pivotal characteristic central to the success of our




Fig. 1. We represent the functional mapping from the density profile $\rho(x)$ to local values of the one-body direct correlation function $c_1(x)$ in planar geometry via a neural network. The density profile is discretized on a regular spatial grid with resolution 0.01σ and given within a region around the location of interest to the input layer. Three fully connected layers with continuously differentiable activation functions enable the inference of the nonlinear and nonlocal functional map. The output layer consists of a single node which yields the predicted value of $c_1(x)$ at the chosen location.

method. To contrast this, assume that one attempts to model the functional mapping $\mu_{\text{loc}}(x) = \mu - V_{\text{ext}}(x) \rightarrow \rho(x)$, thereby naively imitating the simulation procedure. This task poses major challenges due to the long-range nature of density correlations induced by an external potential, which is circumvented in our case by the choice of a more manageable target functional.

The input layer involves 513 nodes and is followed by three fully connected hidden layers with 512 units each. The output layer consists of a single node for the scalar value of $c_1(x)$ at the specified location x . In order to realize a nonlinear input–output mapping of the neural network, activation functions are applied to the output of each node within a hidden layer (see also ref. 43 for a pedagogical introduction to the design of neural networks). We deviate here from the most common choice of a rectified linear unit (ReLU) and instead employ continuously differentiable activation functions such as the exponential linear unit or the softplus function (45). This choice leads to substantial improvements during training and in particular when using automatic differentiation to evaluate two-body quantities; see Sections B.1 and B.2. We attribute the superior performance to the fact that activation functions which are not continuously differentiable and which vanish in certain domain ranges (such as ReLU) reinforce sparsity of the activation output, i.e., the tendency to set many units of a hidden layer identically to zero (46). While this property is desired in many machine learning tasks (e.g., for classification), it hinders the accurate representation of the functional relation $c_1(x; [\rho])$ in our case. The resulting neural functional for the one-body direct correlation function is denoted in the following by $c_1^*(x; [\rho])$ and related quantities which follow from it by inference are marked accordingly by a superscript star.

A.4. Training procedure and metrics. The machine learning routines are implemented in Keras/Tensorflow (43) and we use the standard Adam (47) optimizer for the adjustment of the network parameters in order to fit $c_1^*(x; [\rho])$ against the simulation reference $c_1(x)$. The problem at hand is a regression task. Hence, the mean squared error is chosen as a suitable loss function and the mean average error serves as a validation metric. Since the model shall infer the pointwise value $c_1(x)$ from a density section around a specified location x , see Fig. 1, the simulation data cannot be passed as they are to the neural network. Instead, windowed views of the density profile have to be generated prior to the training loop, which correspond to the target value $c_1(x)$ at the center x of the respective window. A periodic continuation of all simulation profiles is valid due to periodic boundary conditions. Additionally, we use data augmentation to benefit from the inherent mirror symmetry (i.e., $x \rightarrow -x$) of the problem and thus effectively double the number of training datasets. As is customary, we separate the independent simulation results prior to performing the machine learning routines: 150 are kept aside as a test set, 150 serve as validation data to monitor training progress and 450 are used for the actual training of the neural network.

Modeling the functional relationship of $c_1(x; [\rho])$ locally, i.e., inferring pointwise values individually instead of outputting the entire profile at once, has numerous conceptual and practical advantages. Regarding the feasibility of the neural network in concrete applications, one is free to choose an arbitrary box length L when gathering training data and more importantly to readjust the value of L when using the trained neural network for making predictions (cf. Section C.3). From a physical point of view, providing only local density information has the merit of already capturing the correlated short-range behavior of $c_1(x; [\rho])$. If the neural network were to output the entire one-body direct correlation profile from a given density profile $\rho(x)$ at once,

this inherent locality would have to be learned instead, hence leading to a much more elaborate training process. Last, the fine-grained nature of the training data turns out to be highly beneficial from a machine learning perspective. Note that one can generate $9 \cdot 10^5$ input–output pairs from 450 training simulations in the present context (with the values being doubled after data augmentation). The increased cardinality of the training set enables better generalization of the model and also prevents overfitting, e.g., to the statistical noise of the sampled profiles.

We train the model for 100 epochs in batches of size 256 and decrease the learning rate exponentially by $\sim 5\%$ per epoch from an initial value of 0.001. This results in a best mean average error of 0.0022 over the validation set, which is of the same order as the estimated average noise of the simulation data for $c_1(x)$. Therefore, we deem our neural network to possess full representational power of the local functional relationship $c_1(x; [\rho])$ within the conditions of the provided simulation data.

B. Examining the Neural Correlation Functional.

B.1. Two-body bulk correlations. Besides monitoring standard metrics such as the mean average error over a test set, arguably deeper physical insights into the rigorous structure of the statistical mechanics at hand serve for assessing the quality of the neural functional $c_1^*(x; [\rho])$. We first ascertain that the model gives an accurate representation of the physics of bulk fluids. Despite the apparent simplicity of this case, this is a highly nontrivial test as the training data solely covered (strongly) inhomogeneous situations. For this, we investigate the pair structure and aim at implementing the two-body direct correlation functional, which is formally defined as the functional derivative (4)

$$c_2(\mathbf{r}, \mathbf{r}'; [\rho]) = \frac{\delta c_1(\mathbf{r}; [\rho])}{\delta \rho(\mathbf{r}')} \quad [4]$$

On the basis of the neural network, we can make use of the powerful automatic differentiation techniques. This allows to create an immediate analog of Eq. 4 via $c_2^*(x, x'; [\rho]) = \delta c_1^*(x; [\rho]) / \delta \rho(x')$, where the functional derivative $\delta / \delta \rho(x')$ is evaluated by reverse mode automatic differentiation with respect to the input values of the discretized density profile. In common machine learning frameworks, this requires only high-level code, e.g., GradientTape in Keras/Tensorflow (43). The numerical evaluation of $c_2^*(x, x'; [\rho])$ is performant as reverse mode automatic differentiation generates executable code that is suitable for building derivatives with respect to multiple input variables simultaneously.

We obtain the bulk direct correlation function in planar geometry as the special case $\bar{c}_2^b(x, \rho_b) = c_2(0, x; [\rho_b])$, where we have introduced the bulk density $\rho_b(x) = \rho_b = \text{const.}$ (In the notation, the parametric dependence on ρ_b is dropped in the following). Note that $\bar{c}_2^b(x)$ is distinct from the more common radial representation $c_2^b(r)$, as our geometry implies an integration over the lateral directions y and z , i.e.,

$$\begin{aligned} \bar{c}_2^b(x) &= \int dy dz c_2^b\left(r = \sqrt{x^2 + y^2 + z^2}\right) \\ &= 2\pi \int_x^\infty dr r c_2^b(r), \end{aligned} \quad [5]$$

where the last equality follows from using radial coordinates and substitution. The standard radial form $c_2^b(r)$ can however be recovered by differentiating Eq. 5 with respect to x such that

$$c_2^b(r) = -\frac{\tilde{c}_2^{b'}(r)}{2\pi r}, \quad [6]$$

where $\tilde{c}_2^{b'}(r)$ denotes the derivative of $\tilde{c}_2^b(x)$ evaluated at $x = r$. Numerical artifacts might occur particularly for small values of r as evaluating Eq. 6 requires the numerical derivative of $\tilde{c}_2^b(x)$ as well as a division by r .

We perform a Fourier transform of the planar real space representation $\tilde{c}_2^b(x)$ and utilize radial symmetry in Fourier space. This acts as a deconvolution of Eq. 5 and directly yields the radial Fourier (Hankel) transform of $c_2^b(r)$,

$$\tilde{c}_2^b(k) = \frac{4\pi}{k} \int_0^\infty dr r \sin(kr) c_2^b(r). \quad [7]$$

The inverse transform is identical to Eq. 7 up to a factor of $(2\pi)^{-3}$ upon interchanging r and k . To go further, the bulk Ornstein–Zernike equation (4)

$$\tilde{c}_2^b(k) = \frac{\tilde{h}(k)}{1 + \rho_b \tilde{h}(k)}, \quad [8]$$

is used to obtain the total correlation function $\tilde{h}(k)$ from $\tilde{c}_2^b(k)$ in Fourier space after rearrangement. Recall that the radial distribution function follows directly via $g(r) = h(r) + 1$; here, $h(r)$ is the real space representation of $\tilde{h}(k)$. The static structure factor $S(k)$ is then given as:

$$S(k) = 1 + \rho_b \tilde{h}(k). \quad [9]$$

In Fig. 2, results of $\tilde{c}_2^b(x)$, $c_2^b(r)$, $\tilde{c}_2^b(k)$, $\tilde{h}(k)$, and $S(k)$ are shown for different bulk densities $\rho_b \sigma^3 = 0.4, 0.7, 0.9$. From our neural functional, we obtain $\tilde{c}_2^{b*}(x) = \delta c_1^*(0; [\rho]) / \delta \rho(x)|_{\rho=\rho_b}$, i.e. the autodifferentiated network is evaluated at spatially constant density ρ_b . The total correlation function and the static structure factor follow from Eqs. 8 and 9 after having computed $\tilde{c}_2^{b*}(k)$ via a numerical Fourier transform of $\tilde{c}_2^{b*}(x)$. For comparison, we also depict reference data obtained analytically from the Percus–Yevick theory (17) and reproduced from simulation results of Groot et al. (44). Good agreement is found between simulation and the autodifferentiated neural network, while the Percus–Yevick result shows noticeable deviations in $\tilde{c}_2^b(x)$. The latter overestimates the depth of the core region $x < \sigma$, and this discrepancy increases for larger bulk densities. The neural functional yields a clear improvement over the Percus–Yevick theory and shows only marginal differences to the simulation results of ref. 44 for both the planar real space and the radial Fourier space representation of the two-body direct correlation function. In $\tilde{h}(k)$ and $S(k)$, the severity of the discrepancies of simulation and machine learning data to the Percus–Yevick results decreases, but a difference is still noticeable in particular for large bulk densities. A slight mismatch to the simulation reference is observed in the magnitude and phase of the oscillations of the Percus–Yevick static structure factor $S_{PY}(k)$, and this correction is reproduced very well by the neural functional. Note that although one arrives at radial representations of the quantities $\tilde{h}(k)$ and $S(k)$ in Fourier space, performing the radial backtransform to real space numerically according to the inverse of Eq. 7 is generally a “notoriously difficult task” (48) and is not considered here.

This successful test reveals that, while being trained solely with one-body profiles, the neural functional $c_1^*(x; [\rho])$ contains full two-body information equivalent in bulk to the radial

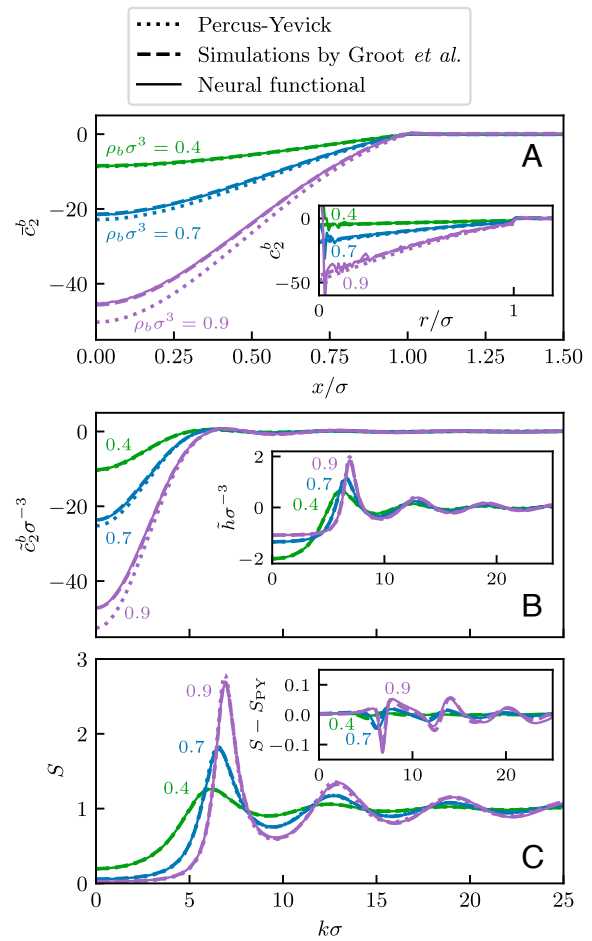


Fig. 2. We compare (A) the planar direct correlation function $\tilde{c}_2^b(x)$, (B) its radial Fourier space representation $\tilde{c}_2^b(k)$, and (C) the static structure factor $S(k)$ for different bulk densities $\rho_b \sigma^3 = 0.4, 0.7, 0.9$ (as indicated). Data are shown as obtained from the Percus–Yevick theory (dotted), from simulation results by Groot et al. (44) and from our neural functional $c_1^*(x; [\rho])$ (solid), where $\tilde{c}_2^{b*}(x)$ is acquired via automatic differentiation. The inset in panel (A) shows the radial direct correlation function $c_2^b(r)$ as obtained via Eq. 6. In panel (B), the inset depicts the total correlation function $\tilde{h}(k)$ in Fourier space, which follows from $\tilde{c}_2^b(k)$ via the bulk Ornstein–Zernike Eq. 8. The inset in panel (C) displays the deviation of $S(k)$ to the Percus–Yevick result $S_{PY}(k)$ for the simulation data and the neural functional. Simulation and machine learning results are in very good agreement with each other while the Percus–Yevick theory shows quantitative discrepancies.

distribution function $g(r)$. The pair correlations can be accessed via automatic differentiation at low computational cost, and they are consistent with known bulk results. We recall that this is a mere by-product of the neural network and that no such two-body information has been explicitly incorporated into the training. More so, Fig. 2 demonstrates that the bulk quantities $\tilde{c}_2^b(x)$, $\tilde{c}_2^b(k)$, $\tilde{h}(k)$, and $S(k)$ as obtained from $c_1^*(x; [\rho])$ substantially outperform the Percus–Yevick theory and almost attain simulation quality. In *SI Appendix*, we illustrate that higher-order correlations such as the three-body direct correlation functional $c_3^*(x, x', x''; [\rho])$ follow analogously via nested automatic differentiation. On this level, differences to FMT results are even more prominent than the deviations to the two-body Percus–Yevick results. As we will show in Section C.2, the accuracy of predictions from the neural network also holds in inhomogeneous situations, where FMT serves again as an analogous and arguably even more challenging theoretical

baseline than the Percus–Yevick bulk theory. Before doing so, we lay out additional consistency tests and quality assessments that are applicable in inhomogeneous systems.

B.2. Noether sum rules. In order to further elucidate whether $c_1^*(x; [\rho])$ quantitatively reproduces fundamental properties of equilibrium many-body systems, we make use of exact sum rules that follow from thermal Noether invariance (42):

$$\nabla c_1(\mathbf{r}) = \int d\mathbf{r}' c_2(\mathbf{r}, \mathbf{r}') \nabla' \rho(\mathbf{r}'), \quad [10]$$

$$\int d\mathbf{r} \rho(\mathbf{r}) \int d\mathbf{r}' \rho(\mathbf{r}') \nabla c_2(\mathbf{r}, \mathbf{r}') = 0. \quad [11]$$

Both Eqs. 10 and 11 apply in any equilibrated inhomogeneous system regardless of the type of internal interactions. While the interparticle interaction potential does not appear explicitly in Eqs. 10 and 11, it nevertheless determines the functionals $c_1(\mathbf{r}; [\rho])$ and $c_2(\mathbf{r}, \mathbf{r}'; [\rho])$. Recall that the spatial gradient of the one-body direct correlation function can be identified with the internal equilibrium force profile, $\mathbf{f}_{\text{int}}(\mathbf{r}) = k_B T \nabla c_1(\mathbf{r})$ (41).

We verify that the neural functional complies with the above sum rules Eqs. 10 and 11 as follows. Analogous to Section B.1, we use autodifferentiation to evaluate Eq. 4, but this time, we retain the full inhomogeneous structure of $c_2^*(x, x'; [\rho])$. The left-hand side of Eq. 10 is obtained straightforwardly from simple evaluation of the neural functional and numerical spatial differentiation. As input for $\rho(x)$, we use the simulated density profiles of the test set. Care is required when evaluating the spatial gradients $\nabla \rho(x)$, $\nabla c_1^*(x; [\rho])$ and $\nabla c_2^*(x, x'; [\rho])$ due to the amplification of undesired noise, which we reduce by applying a low-pass filter after having taken the numerical derivatives. The volume integrals reduce in planar geometry to $\int d\mathbf{r} = A \int dx$, where A is the lateral system area.

In Fig. 3, three typical profiles for the left- and right-hand side of Eq. 10 are shown. In all three systems, both sides of the equation coincide up to numerical noise due to the required spatial derivatives. Additionally, we define errors via scalar deviations from equality in Eqs. 10 and 11 respectively as

$$e_1 = \left\| \nabla c_1(x) - A \int dx' c_2(x, x') \nabla' \rho(x') \right\|_{\infty}, \quad [12]$$

$$e_2 = A^2 \int dx \rho(x) \int dx' \rho(x') \nabla c_2(x, x'), \quad [13]$$

where $\|\cdot\|_{\infty}$ denotes the maximum norm. Fig. 3 A and B depict results for e_1 and e_2 resolved with respect to the mean density $\bar{\rho} = \int d\mathbf{r} \rho(\mathbf{r})/V$ for all 150 density profiles of the test set, where V denotes the volume of the system. The small magnitudes of the observed error values indicate that the neural network satisfies the Noether identities 10 and 11 to very high accuracy. Outliers are caused mostly by the moderate numerical noise of the spatial gradients (Fig. 3, III) and are no hindrance in practical applications of the neural functional.

This confirmation demonstrates that our method transcends the neural network from a mere interpolation device of the simulation training data to a credible standalone theoretical object. The fact that one is able to carry out consistent and performant functional calculus indeed renders $c_1^*(x; [\rho])$ a neural-network-based density functional. Besides functional differentiation, we show next that functional line integration acts as the inverse operation and provides access to the corresponding free energy. SI Appendix gives further insight into the symmetry properties of $c_2^*(x, x'; [\rho])$, which serve as a prerequisite for the




Fig. 3. Typical profiles of the right- and left-hand sides of Eq. 10 are shown for three test scenarios in panels (I), (II), and (III), where one can verify their high level of agreement across the entire inhomogeneous systems. Additionally, the respective scalar discrepancies e_1 and e_2 of the Noether identities Eqs. 10 and 11 are displayed, defined as (A) the maximum norm of the difference of left- and right-hand side of Eq. 10 and (B) the value of the left-hand side of Eq. 11. Across all mean densities $\bar{\rho}$ of the test set, the sum rules are satisfied to very high accuracy by our model. Some outliers remain which we attribute to the numerical computation of spatial gradients appearing in Eqs. 10 and 11; see also panel (III) for an example of the noise that this introduces in the respective terms of Eq. 10 particularly in the vicinity of hard walls.

existence of a generating excess free energy functional $F_{\text{exc}}^*[\rho]$; we recall Eq. 2.

B.3. Equation of state and free energy. Although the machine learning procedure operates on the level of the one-body direct correlation function, the excess free energy $F_{\text{exc}}[\rho]$ is accessible by functional line integration (49):

$$\beta F_{\text{exc}}[\rho] = - \int_0^1 d\alpha \int d\mathbf{r} \rho(\mathbf{r}) c_1(\mathbf{r}; [\rho_\alpha]). \quad [14]$$

Here, $\rho_\alpha(\mathbf{r}) = \alpha \rho(\mathbf{r})$ is a sequence of density profiles that are linearly parametrized by α in the range $0 \leq \alpha \leq 1$. The limits are $\rho_0(\mathbf{r}) = 0$ such that $F_{\text{exc}}[0] = 0$, and $\rho_1(\mathbf{r}) = \rho(\mathbf{r})$, which is the target density profile that appears as the functional argument on the left-hand side of Eq. 14. Other parametrizations of $\rho_\alpha(\mathbf{r})$ are conceivable but change the concrete form of Eq. 14. On the basis of $c_1^*(x; [\rho])$, we implement Eq. 14 via $\beta F_{\text{exc}}^*[\rho] = -A \int_0^1 d\alpha \int dx \rho(x) c_1^*(x; [\rho_\alpha])$ and evaluate the integrals numerically; as before, A denotes the lateral system area.

We first return to bulk systems and illustrate in the following three different routes toward obtaining the bulk equation of state from the neural network. For this, we introduce the excess free energy density as $\psi_b(\rho_b) = F_{\text{exc}}[\rho_b]/V$. From the neural functional, the excess free energy density $\psi_b^*(\rho_b)$ can be acquired via $F_{\text{exc}}^*[\rho_b]$ from functional line integration along a path of bulk densities according to Eq. 14. Alternatively and equivalently, one can simply evaluate the neural direct correlation functional at bulk density ρ_b and due to translational symmetry at arbitrary location (e.g., $x = 0$) such that $c_1^{b*} = c_1^*(0; [\rho_b])$. Simplifying Eq. 2 in bulk reveals that

$$\psi_b^{*'}(\rho_b) = -k_B T c_1^{b*}, \quad [15]$$

where the prime denotes the derivative with respect to the bulk density argument. The excess free energy density $\psi_b^*(\rho_b)$ follows

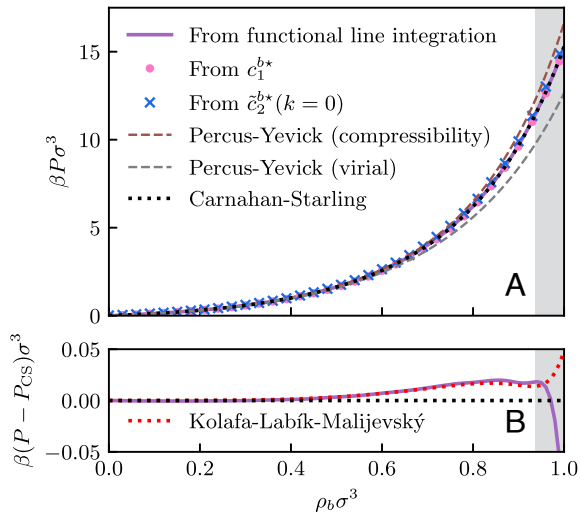


Fig. 4. We show (A) the equation of state $P(\rho_b)$ obtained via different methods and (B) deviations to the Carnahan–Starling result $P_{CS}(\rho_b)$ (dotted black line). The neural equation of state $P^*(\rho_b)$ is calculated via Eq. 16 in which the excess free energy density follows from functional line integration according to Eq. 14 (solid purple line), from evaluation of the bulk value c_1^{b*} (pink dots), see Eq. 15, and via the low-wavelength limit of $c_2^{b*}(k)$ (blue crosses), see Eq. 17. For comparison, the Percus–Yevick equations of state according to the virial (dashed gray line) and compressibility (dashed brown line) route are shown. Bulk densities ρ_b beyond the stable fluid phase are shaded in gray. All three routes coincide very well up to and within the metastable region, with functional line integration leading to the most accurate results. We additionally depict a simulation-based equation of state (B, dotted red line) due to Kolafa et al. (50), which our neural functional is able to reproduce very accurately in the stable fluid region, hence exceeding in precision the Carnahan–Starling equation of state.

from ordinary numerical integration across bulk densities up to the target value ρ_b . The numerical accuracy to which both routes coincide serves as a further valuable consistency test.

Additionally, one obtains the bulk pressure $P(\rho_b)$ from the excess free energy density via

$$P(\rho_b) = (\psi_b'(\rho_b) + k_B T) \rho_b - \psi_b(\rho_b). \quad [16]$$

The pressure is equally accessible from a further route which incorporates previous results for the bulk pair structure via their low-wavelength limits according to ref. (4):

$$\beta \left. \frac{\partial P}{\partial \rho_b} \right|_T = \frac{\beta}{\rho_b \chi_T} = \frac{1}{S(0)} = \frac{1}{1 + \rho_b \tilde{b}(0)} = 1 - \rho_b \tilde{c}_2^b(0), \quad [17]$$

where one can identify the isothermal compressibility $\chi_T = \rho_b^{-1}(\partial \rho_b / \partial P)_T$. From Eq. 17, $P(\rho_b)$ is obtained by evaluation of either of the bulk correlation functions (Section B.1) in Fourier space at $k = 0$ for different bulk densities and by subsequent numerical integration toward the target value of ρ_b .

We compare the results in Fig. 4, where the equation of state $P^*(\rho_b)$ of the neural network was acquired from functional line integration across bulk systems, cf. Eq. 14, from evaluation of one-body bulk correlation values c_1^{b*} , cf. Eq. 15, and from the low-wavelength limit of two-body bulk correlations, cf. Eq. 17. One finds that the results of all three routes are consistent with each other and that they match very well the Carnahan–Starling equation of state (51), thus outperforming the Percus–Yevick theory as already observed for the bulk pair structure in Section B.1. A slight deviation can be noticed when evaluating $P^*(\rho_b)$ via Eq. 17, which constitutes the

most indirect route detouring to two-body correlations. This may reflect the small discrepancy of the neural functional to simulation results (cf. Fig. 2) and the sensitivity of the low-wavelength limit of the static structure factor to remaining finite size effects (52). Notably, functional line integration is the most reliable method, and the corresponding results even surpass the Carnahan–Starling equation of state in accuracy. Fig. 4B shows the reproduction of a highly accurate simulation-based equation of state due to Kolafa et al. (50). We recall again that neither bulk information nor data for free energies or pressures was given explicitly in the training of the neural network. Instead, the beyond-Carnahan–Starling precision is achieved solely by extracting direct one-body correlations from simulation data of randomized inhomogeneous systems in planar geometry. In *SI Appendix*, we additionally demonstrate that the neural functional is fit for the application of dimensional crossover (53) in order to obtain the bulk equation of state for the two-dimensional hard disk fluid within a reasonable range of packing fractions.

For a concise comparison of free energies in inhomogeneous situations, additional reference data have to be acquired from simulations. In our grand canonical setting, thermodynamic integration (54) with respect to the chemical potential can be used to measure the grand potential according to

$$\Omega[\rho] = - \int_{-\infty}^{\mu} d\mu' \langle N \rangle. \quad [18]$$

Here, the integration starts from an empty system with $\Omega[0] = 0$ and traverses the chemical potential up to the target value μ . One needs to measure the mean number of particles $\langle N \rangle$ in a sufficient number of simulations with intermediate chemical potentials $-\infty < \mu' \leq \mu$ to evaluate Eq. 18 numerically. The excess free energy then follows directly from

$$F_{\text{exc}}[\rho] = \Omega[\rho] - F_{\text{id}}[\rho] - \int d\mathbf{r} \rho(\mathbf{r})(V_{\text{ext}}(\mathbf{r}) - \mu), \quad [19]$$

where $F_{\text{id}}[\rho] = k_B T \int d\mathbf{r} \rho(\mathbf{r})(\ln \rho(\mathbf{r}) - 1)$ is the ideal gas free energy. Thermodynamic integration according to Eq. 18 has been performed for 22 systems of the test set to yield reference values $F_{\text{exc}}^{\text{sim}}$ for the excess free energy via Eq. 19. The systems were selected to cover a broad range of excess free energy values, and FMT results for F_{exc} were used as a further theoretical estimate for this selection.

In Table 1 and Fig. 5, we show errors of F_{exc} to the quasi-exact simulation values when calculating the excess free energy via Rosenfeld and White Bear MkII FMT as well as from functional

Table 1. The absolute and relative mean average error of the excess free energy F_{exc} as obtained via the Rosenfeld and White Bear MkII FMT functionals is compared to the result from functional line integration of the neural correlation functional

	$\beta \langle F_{\text{exc}} - F_{\text{exc}}^{\text{sim}} \rangle$	$\langle F_{\text{exc}} - F_{\text{exc}}^{\text{sim}} / F_{\text{exc}}^{\text{sim}} \rangle$
Rosenfeld	0.540	1.75%
White Bear MkII	0.0159	0.104%
Neural functional	0.0127	0.097%

The reference values $F_{\text{exc}}^{\text{sim}}$ were obtained via thermodynamic integration according to Eqs. 18 and 19 for a subset of the test systems. The results of the neural functional surpass the Rosenfeld FMT significantly and even yield a slight improvement over the highly accurate White Bear MkII theory. The angular brackets denote an average over the 22 test simulations.




Fig. 5. We compare free energies of inhomogeneous test systems as obtained via Rosenfeld (turquoise squares) and White Bear (purple triangles) FMT as well as with functional line integration of the neural correlation functional $c_1^*(x; [\rho])$ (yellow circles). The discrepancy $F_{\text{exc}} - F_{\text{exc}}^{\text{sim}}$ of the respective method to the simulation result $F_{\text{exc}}^{\text{sim}}$ is shown. Rosenfeld FMT systematically underestimates F_{exc} , whereas White Bear MkII FMT as well as our neural functional yield almost exact results. The neural network performs slightly better for large excess free energies as occur primarily in dense systems.

line integration according to Eq. 14 of the neural functional. For both FMT methods, a DFT minimization (cf. Section C.1) is performed to yield a self-consistent density profile $\rho(x)$, which serves as input to the respective analytic FMT expression for $F_{\text{exc}}[\rho]$. Hence, we compare consistently equilibrium states (according to the respective theory) corresponding to the same form of the external potential.

The comparison reveals that the neural functional significantly outperforms Rosenfeld FMT and still yields slightly more accurate values for the excess free energy than the very reliable White Bear theory. Regarding the above-described bulk results for the free energy, this behavior is both consistent and expected, as the Rosenfeld and White Bear MkII functionals can be associated with the Percus–Yevick compressibility and Carnahan–Starling bulk equations of state respectively. Still, the test in inhomogeneous systems is a more rigorous one than in bulk, as the full nonlocal functional representation is invoked when providing $c_1^*(x; [\rho])$ with an inhomogeneous density profile as input. Given that the functional line integration of $c_1^*(x; [\rho])$ via Eq. 14 is practically immediate, one can deem $F_{\text{exc}}^*[\rho]$ itself a corresponding neural functional for the excess free energy that enables a full description of the thermodynamics of inhomogeneous fluids to high accuracy. As we present below, this quantitative precision is preserved when applying the neural functional in a predictive manner in the self-consistent calculation of density profiles.

C. Predicting Inhomogeneous Fluids via Neural DFT.

C.1. Going beyond analytic approximations. In the previous section, the trained model has been put to test by deriving related quantities such as $c_2^*(x, x'; [\rho])$ from autodifferentiation and $F_{\text{exc}}^*[\rho]$ from functional line integration in order to assess its performance against analytic and numerical reference results. We now turn to the application of the neural functional $c_1^*(x; [\rho])$ in the context of the self-consistent determination of density profiles according to the DFT Euler–Lagrange equation. This is achieved by rearranging Eq. 1 to the standard form (2, 4)

$$\rho(\mathbf{r}) = \exp(-\beta(V_{\text{ext}}(\mathbf{r}) - \mu) + c_1(\mathbf{r}; [\rho])). \quad [20]$$

A fixed-point (Picard) iteration with mixing parameter α can be used to determine the density profile from Eq. 20 according to

$$\rho(\mathbf{r}) \leftarrow (1 - \alpha)\rho(\mathbf{r}) + \alpha \exp(-\beta(V_{\text{ext}}(\mathbf{r}) - \mu) + c_1(\mathbf{r}; [\rho])). \quad [21]$$

The degree of convergence is determined from the remaining difference of right- and left-hand side of Eq. 20. With the trained neural functional at hand, one can evaluate the one-body direct correlation function in Eq. 21 via the surrogate $c_1^*(x; [\rho])$ in each iteration step. In the following, the use of $c_1^*(x; [\rho])$ in this context will be referred to as neural DFT.

We note two minor technical points concerning the use of the neural functional in the Picard iteration. It was observed that a conservative choice of α is necessary during the first few iterations to ensure numerical stability. After this burn-in, the mixing parameter can be set to usual values (e.g. $\alpha = 0.05$). Furthermore, the convergence criterion has to be relaxed as compared to typical choices in analytic DFT methods due to the remaining intrinsic uncertainty of $c_1^*(x; [\rho])$. The mean average error after training, cf. Section A.4, provides an estimate for the expected relative uncertainty of the density profile according to Eq. 20. Depending on the specific problem, the error might not decrease any further than that during the iteration of Eq. 21. Neither of these points caused any practical hindrance in applications.

The treatment of Eq. 20 in neural DFT is conceptually not different from that in standard DFT methods. However, the model $c_1^*(x; [\rho])$ relieves the theory from being restricted by the available approximations for the one-body direct correlation function as generated from analytic expressions of the excess free energy functional $F_{\text{exc}}[\rho]$ via Eq. 2. We emphasize that, unlike in previous work (35, 37), no analytic ansatz had to be provided and that our method is generic for the determination of a suitable functional from a given model Hamiltonian, thus indeed constituting a “machine learning black box” (35) regarding the training procedure. However, in contrast to a closed black box, the inner workings of the resulting neural correlation functional can be inspected very thoroughly via the neural functional calculus laid out above. Also note that, while the model works at the level of the one-body direct correlation function, the free energy is readily available from functional line integration, cf. Section B.3. Last, we point out that $c_1^*(x; [\rho])$ captures the entirety of the intrinsic correlations and that further improvements are conceivable by only learning differences to an analytic reference functional. To demonstrate the capabilities of our method, we refrain from this route and show that the trained neural functional alone already exceeds the accuracy of FMT.

C.2. Comparison to FMT. In the following, we benchmark the self-consistent inhomogeneous density profiles obtained via neural DFT against FMT results. For this comparison, the Rosenfeld (8) and White Bear MkII (13) FMT functionals are considered and the simulated density profiles are taken as quasi-exact reference data. The FMT functionals are the most profound analytic description of the hard sphere fluid with the White Bear MkII theory being the state-of-the-art treatment of short-ranged intermolecular repulsion in classical DFT. Nevertheless, measurable and systematic deficiencies still remain, e.g., in highly correlated systems (55). We point the reader to ref. 14 for a thorough account of FMT and to ref. 56 for a very recent quantitative assessment. Note that the tensorial weights of Tarazona (15) to describe hard sphere freezing are not included in our investigation.

The comparison is set up as follows. For each hard sphere system of the test set (Section A.4), we determine the density profile $\rho(x)$ from the Rosenfeld and White Bear MkII FMT functionals as well as from $c_1^*(x; [\rho])$ via the Picard iteration Eq. 21 of the Euler–Lagrange Eq. 20. For this, only the known

form of the external potential $V_{\text{ext}}(x)$ and the value μ of the chemical potential are prescribed. As reference density profiles are available from GCMC simulations, we can evaluate the error $\Delta\rho(x)$ of each of the DFT results relative to the simulation data for $\rho(x)$. From here, different scalar metrics for the quantitative agreement of self-consistent DFT profiles and simulation results are considered.

In Fig. 6, both global and local error measures for the deviation of FMT as well as neural DFT to simulation data are depicted. For the assessment of the global error, we show the L_2 -norm $\|\Delta\rho\|_2$ of the discrepancy to the reference profile, which is normalized by the mean density $\bar{\rho}$ of each system respectively. As the test data cover very dilute to very dense systems, this relative global error measure is plotted as a function of $\bar{\rho}$ to discern the behavior with respect to varying global average density. Similarly, we define an estimate for the relative local error by evaluating the maximum norm $\|\Delta\rho\|_\infty$ of the density deviation divided by the maximum value $\|\rho\|_\infty$ of the GCMC density profile. This quantity is resolved against the maximum $\|\rho\|_\infty$ of the respective

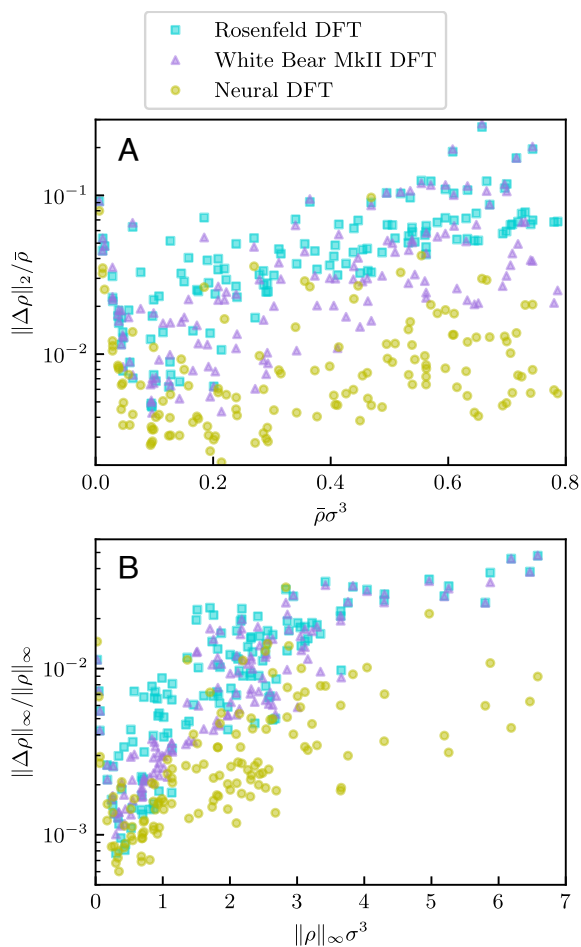


Fig. 6. Measures of discrepancy of self-consistent density profiles to simulation results across the test set are presented. We show (A) the normalized L_2 -norm $\|\Delta\rho\|_2/\bar{\rho}$ as a function of the mean density $\bar{\rho}$ for judgment of the average error over the inhomogeneous system, and (B) the relative maximum norm $\|\Delta\rho\|_\infty/\|\rho\|_\infty$ as a function of the largest local density $\|\rho\|_\infty$ to reveal the magnitude of local errors, e.g. at density peaks. The self-consistent density profiles are obtained from Rosenfeld (turquoise squares) and White Bear MkII (purple triangles) FMT (8, 13) as well as from employing our neural functional $c_1^*(x; [\rho])$ in the DFT Euler–Lagrange equation (yellow circles). Regarding both global and local error, the neural network outperforms the analytic FMT functionals and reduces the respective errors up to an order of magnitude, especially in large-density regimes.

inhomogeneous density, thus enabling the detection of local discrepancies, e.g., in the vicinity of maxima and discontinuities of the density profile.

One recognizes that neural DFT yields substantially better results than the FMT functionals with regard to both error measures. Compared to the Rosenfeld results, both the global and the local errors are decreased by approximately an order of magnitude. Surprisingly, even the White-Bear MkII functional is not able to match the accuracy of the neural DFT, which is noticeable especially for large values of $\bar{\rho}$ and of $\|\rho\|_\infty$.

C.3. Simulation beyond the box. A particular advantage of the local nature of the neural functional $c_1^*(x; [\rho])$ is its applicability to systems of virtually arbitrary size. As explained in Section A.3, it is sufficient to provide the density profile within a rather narrow window as input to the neural network to infer the value of the one-body direct correlation function at the center of the density section. The model $c_1^*(x; [\rho])$ can therefore be used directly in the Euler–Lagrange Eq. 20 for the prediction of planar systems of arbitrary length. Due to the low computational demands of solving this equation self-consistently, this method is suitable even in multiscale problems where macroscopic length scales compete with and are influenced by microscopic correlations and packing features. Although one could argue that analytic DFT methods already account for such tasks, importantly the neural functional $c_1^*(x; [\rho])$ acts as a drop-in replica of the (almost) simulation-like description of the intrinsic correlations. Therefore, neural DFT facilitates to fuse simulation data with common DFT methods, thus providing a means to “simulate beyond the box.”

Simulation beyond the box is demonstrated in Fig. 7, where the system size has been increased to $1,000\sigma$ while the numerical grid size remains unchanged at 0.01σ . Our setup implies that for colloids of, say, size $\sigma = 1\ \mu\text{m}$, we have spatial resolution of 10 nm across the entirety of a system of macroscopic size 1 mm. We consider both a highly correlated fluid in a rapidly varying external potential as well as the diffusive sedimentation behavior (57) in a weak gravitational potential. The former case is realized by generating a sequence of randomized external potentials via Eq. 3 which are spatially connected; the chemical potential is set to zero. Neural DFT yields a highly inhomogeneous density profile in this system and resolves the microscopic variations accurately at low computational cost. In the sedimentation column, a local chemical potential $\mu_{\text{loc}}(x) = \mu - V_{\text{ext}}(x) = (10 - 0.01x/\sigma)k_B T$ is imposed which decreases linearly with respect to the height x , and the system is bounded from the bottom ($x = 0$) and the top ($x = 1,000\sigma$) by hard walls. The spatial variation of $\mu_{\text{loc}}(x)$ is chosen small enough to enable thermal diffusion across the whole sedimentation column and to yield locally an almost bulk-like behavior except near the upper and lower hard walls. The method reproduces both the highly correlated nature of $\rho(x)$ in the vicinity of the walls as well as its intermediate behavior within the sedimentation column, which follows closely the bulk equation of state (Section B.3), as one would expect within a local density approximation (4). In both cases, the computational cost for the determination of $\rho(x)$ with neural DFT is negligible as compared to analogous many-body simulations, which are hardly feasible on such length scales.

Discussion

In this work, we have outlined and validated a machine learning procedure for representing the local functional map from the density profile to the one-body direct correlation function via a neural network. The resulting neural functional was shown

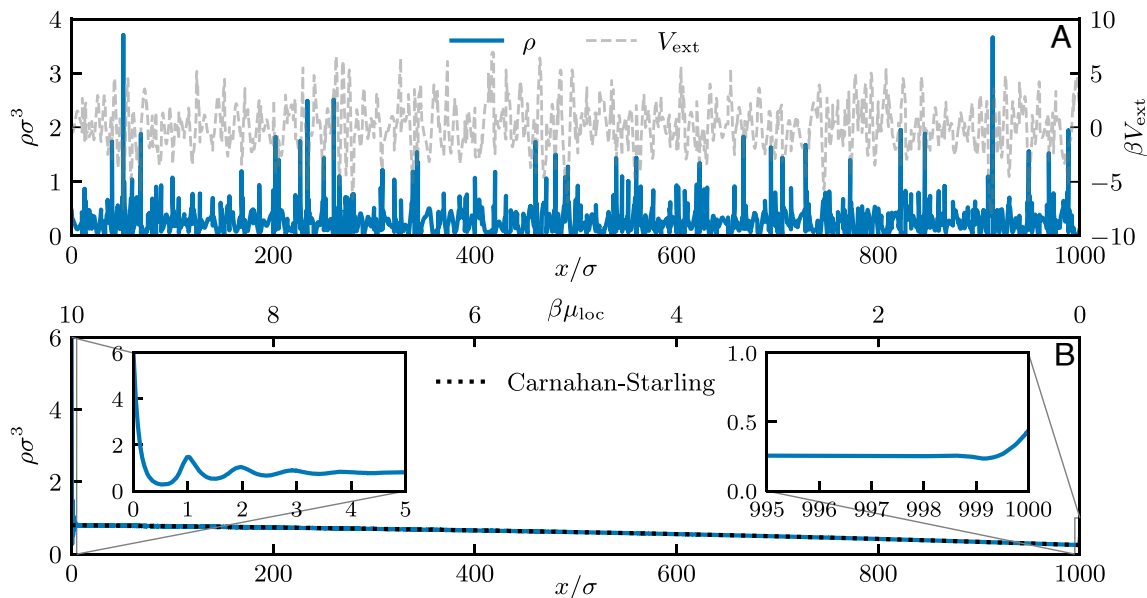


Fig. 7. Neural DFT is used to obtain the density profile $\rho(x)$ (blue lines) of the hard sphere fluid (A) in a highly correlated system with randomized external potential $V_{\text{ext}}(x)$ (gray dashed line) and (B) in a sedimentation column of height $1,000\sigma$ that is bounded by hard walls at the bottom and at the top of the system. Near-simulation microscopic accuracy is retained at low computational cost by the application of neural DFT in the highly correlated large-scale system. For the case of sedimentation, strongly oscillating behavior at the lower wall as well as mild adsorption at the top can be resolved. As the spatial variation of the local chemical potential $\mu_{\text{loc}}(x)$ is negligible, the density profile reproduces the equation of state within the sedimentation column, which is verified by a comparison to the Carnahan–Starling equation of state (dotted black line).

to be applicable as a powerful surrogate in the description of inhomogeneous equilibrium fluids. This was demonstrated for the hard sphere fluid, where we have used GCMC simulations in randomized inhomogeneous planar environments for the generation of training, validation, and test data. Density and one-body direct correlation profiles followed respectively from direct sampling and from evaluation of Eq. 1.

DFT elevates the role of the one-body direct correlation function $c_1(x)$ to that of an intrinsic functional $c_1(x; [\rho])$ depending on the density profile $\rho(x)$ but being independent of the external potential. We exploited this fact in the construction of our neural network, which takes as input a local section of the discretized density profile around a fixed location x and outputs the value of the one-body direct correlation functional $c_1(x; [\rho])$ at that specific location. Establishing a pointwise inference of $c_1(x; [\rho])$ instead of trying to represent the global functional mapping of the entire one-body profiles comes with various advantages, such as independence of the box size, the correct description of the short-range behavior of $c_1(x; [\rho])$, and a very significant improvement of training statistics.

The nonlinear and nonlocal functional relationship was realized by fully connected hidden layers with smooth activation functions and a standard supervised training routine was used. The achieved mean average error over the test set was of the same order of magnitude as the noise floor of the simulations, thus being indicative of full representational power of the neural correlation functional within the considered simulation data. Whether the quality of the model can be improved further by performing more extensive sampling to reduce the statistical noise of the simulation profiles remains to be investigated in the future. Additionally, active and reinforcement machine learning techniques could be useful for interleaving the training and simulation process, thereby guiding the generation of reference data in order to explore the space of inhomogeneous systems more efficiently and exhaustively.

The neural functional was put to test by verifying numerous physical relations in bulk and in inhomogeneous systems.

In particular, it was shown that the two-body direct correlation functional $c_2(x, x'; [\rho])$ as well as higher-order correlations are accessible from the model via automatic differentiation. In bulk, the pair structure as described by the neural network significantly outperforms the Percus–Yevick theory and is even able to compete with simulation results (44), although no bulk data was used during training. In inhomogeneous situations, the conformance of the neural functional to the thermal Noether sum rules Eqs. 10 and 11 as well as to spatial symmetry requirements holds to high accuracy. The excess free energy $F_{\text{exc}}[\rho]$ is readily and efficiently available via functional line integration of the model according to Eq. 14 and the results agree with those obtained from simulations. The bulk equation of state can be acquired consistently from various routes with the results attaining simulation quality (50) and in particular exceeding the very reliable Carnahan–Starling equation of state (51) in accuracy. Dimensional crossover is feasible for the calculation of the bulk equation of state for the two-dimensional hard disk system.

Arguably the most important consequence of the neural functional framework is the applicability of $c_1^*(x; [\rho])$ in the self-consistent calculation of density profiles by solving the Euler–Lagrange equation (20) of classical DFT. As the one-body direct correlation function is faithfully represented by the neural network, one is exempted from having to find analytic approximations for $c_1(x; [\rho])$ or for its generating functional $F_{\text{exc}}[\rho]$. Although FMT provides such approximations for the hard sphere fluid with high precision, we could demonstrate that our neural functional outperforms both the Rosenfeld (8) as well as the White Bear MkII (13) functional. For this, Eq. 20 was solved self-consistently for all 150 randomized local chemical potentials of the test set to obtain $\rho(x)$, where $c_1(x; [\rho])$ was given either analytically by FMT or evaluated via $c_1^*(x; [\rho])$. The comparison of the results to the simulated density profiles reveals that neural DFT yields global and local errors that are up to an order of magnitude lower than those of FMT.

Furthermore, due to the flexibility that comes with the local functional mapping, the neural network could be used as a means

to “simulate beyond the box.” That is, while the training was based solely on simulation data from systems of manageable size, the resulting model $c_1^*(x; [\rho])$ is directly applicable for predictions on much larger length scales. We demonstrated this by imposing a spatial sequence of randomized external potentials on a length of $1,000\sigma$. While the explicit numerical simulation of such a system is comparatively cumbersome, neural DFT offers a way to achieve close to simulation-like accuracy at low computational effort. Furthermore, we have considered a sedimentation column with a height of $1,000\sigma$ that is bounded by hard walls. Neural DFT is capable of both resolving microscopically the adsorption at the walls as well as efficiently capturing the long-range density decay with increasing height. The presented fusion of machine learning and DFT can therefore be another useful technique to make headway in the multiscale description of soft matter (58–60).

Even though we saw no need for a more sophisticated training procedure in our investigations, it could be useful to consider physics-informed machine learning (61) as a technique for enforcing exact physical relations of the underlying problem directly during training. Sum rules in bulk or in inhomogeneous systems, e.g., the thermal Noether identities Eqs. 10 and 11, might be suitable candidates for this task. Analogous to the evaluation of derivatives in physics-informed neural networks, we have shown the necessary quantities to be accessible by automatic differentiation of the neural functional.

When considering nonequilibrium systems, power functional theory (PFT) (40, 41) establishes an exact functional many-body framework which is analogous to that of DFT in equilibrium. A central ramification of PFT is the existence of a functional map from the time-dependent one-body density $\rho(\mathbf{r}, t)$ and current $\mathbf{J}(\mathbf{r}, t)$ to the internal force profile $\mathbf{f}_{\text{int}}(\mathbf{r}, t; [\rho, \mathbf{J}])$, which is in general nonlocal in space and causal in time t . Recent work by de las Heras et al. (39) demonstrated that machine learning this kinematic internal force functional yields highly promising results and overcomes the analytic and conceptual limitations of dynamical density functional theory. In this regard, our method can be put into a more general context as it may be viewed as a mere special case for equilibrium systems where $\mathbf{J}(\mathbf{r}, t) = 0$. The topical problem of accurately describing nonequilibrium many-body physics is certainly a natural contender for the application and extension of our neural functional framework, with many practical questions arising, e.g., concerning the generation of training data or the choice of neural network architecture.

While much insight could be gained by considering the hard sphere fluid, the application of our machine learning procedure is arguably even more useful for particle models that lack satisfactory analytic density functional approximations. Although mean-field descriptions account surprisingly well for soft and attractive contributions (62, 63), e.g., in the Lennard-Jones fluid, analytic efforts to go beyond this approximation are sparse (64–66). We demonstrate the generality of our method in *SI Appendix*, where we show that the machine learning routine applies directly to the (truncated) Lennard-Jones interaction potential in an isothermal supercritical setting. In the future, providing the temperature as a further input quantity to a modified neural network is a valuable goal in order to tackle the full physics of such thermal systems. As a proper treatment of the arising phase transitions and interfacial phenomena is already subtle in simulation, the machine learning perspective might provide further insights. We expect the general method to hold up even for complex particle models, e.g., containing many-body interactions (22), provided that sufficiently accurate training data of sufficient quantity can be generated.

For the treatment of anisotropic particles, the neural network must be extended to accommodate for the additional orientational degrees of freedom. Recent advances in molecular DFT could be helpful in guiding appropriate augmentations of our method (67, 68). Related to the increased dimensionality due to anisotropy, the extension of the machine learning procedure from planar symmetry to more general geometries is worth contemplating. Especially for fully inhomogeneous three-dimensional problems, the amount of required training data seems restrictive at first. However, we have shown in this work that results obtained in planar geometry already capture the essence of internal interactions. Therefore, it may be feasible to base the machine learning predominantly on data in reduced geometrical settings and to incorporate remaining nontrivial effects due to the more general geometry by supplementing only a few selected higher-dimensional simulations. In particular, we highlight in this context the promising development of equivariant neural networks (69–72), which serve as a means of casting underlying symmetries of a problem directly into the neural network architecture. Recent applications in the physical domain show that this method facilitates robust training and generalization on the basis of much reduced datasets as compared to common machine learning approaches which do not intrinsically enforce symmetry (73–75). In our case, exploiting inherent symmetries of the direct correlation functional via the use of equivariant neural networks is certainly valuable when further orientational or spatial degrees of freedom are to be considered.

Last, we point out useful cross-fertilization of machine learning ideas regarding topical applications in quantum DFT (76). In particular, the analogous functional mapping to the classical one-body direct correlation functional $c_1(\mathbf{r}; [\rho])$ is given quantum mechanically by the exchange-correlation potential $v_{\text{xc}}(\mathbf{r}; [n])$ which depends functionally on the one-body electron density $n(\mathbf{r})$. Due to the immediate analogy, obtaining the exchange-correlation energy functional $E_{\text{xc}}[n]$ might be feasible with functional line integration similar to our treatment of $F_{\text{exc}}[\rho]$ via Eq. 14, which here becomes $E_{\text{xc}}[n] = \int_0^1 d\alpha \int d\mathbf{r} n(\mathbf{r}) v_{\text{xc}}(\mathbf{r}; [n_\alpha])$ with $n_\alpha(\mathbf{r}) = \alpha n(\mathbf{r})$. Albeit lacking the neural functional calculus that we presented here, Zhou et al. (77) have successfully demonstrated the machine learning of the functional mapping from the electron density to local values of the exchange-correlation potential $v_{\text{xc}}(\mathbf{r})$. Specifically, they trained a convolutional neural network on the basis of three-dimensional quantum chemical simulation data of small molecules and could obtain accurate predictions for larger molecules. This success is akin to the multiscale applicability of our neural correlation functional $c_1^*(\mathbf{r}; [\rho])$. In general, however, most machine learning strategies in quantum DFT have considered different functional mappings (78–83). In light of our results for classical systems, we deem the analogous machine learning of the local functional relationship of $v_{\text{xc}}(\mathbf{r}; [n])$ the arguably most promising approach in the development of a neural quantum DFT with the goal of chemical accuracy and generic applicability.

Data, Materials, and Software Availability. Code, data sets and models have been deposited in Zenodo (84)

ACKNOWLEDGMENTS. We thank T. Zimmermann, T. Eckert, and N. C. X. Stuhlmueller for useful comments. This work is supported by the German Research Foundation (DFG) via Project No. 436306241.

Author affiliations: ^aTheoretische Physik II, Physikalisches Institut, Universität Bayreuth, Bayreuth D-95447, Germany

1. R. Evans, J. P. Hansen, H. Löwen, Foreword. *J. Phys.: Condens. Matter* **14**, 1 (2002).
2. R. Evans, The nature of the liquid-vapour interface and other topics in the statistical mechanics of non-uniform, classical fluids. *Adv. Phys.* **28**, 143–200 (1979).
3. R. Evans, *Density Functionals in the Theory of Non-Uniform Fluids* D. Henderson, Eds. (M. Dekker, New York, 1992), pp. 85–175.
4. J. P. Hansen, I. R. McDonald, *Theory of Simple Liquids With Applications to Soft Matter* (Elsevier Science & Technology Books, 2013).
5. R. Evans, M. Oettel, R. Roth, G. Kahl, New developments in classical density functional theory. *J. Phys.: Condens. Matter* **28**, 240401 (2016).
6. N. D. Mermin, Thermal properties of the inhomogeneous electron gas. *Phys. Rev.* **137**, A1441–A1443 (1965).
7. J. K. Percus, Equilibrium state of a classical fluid of hard rods in an external field. *J. Statist. Phys.* **15**, 505–511 (1976).
8. Y. Rosenfeld, Free-energy model for the inhomogeneous hard-sphere fluid mixture and density-functional theory of freezing. *Phys. Rev. Lett.* **63**, 980–983 (1989).
9. E. Kierlik, M. L. Rosinberg, Free-energy density functional for the inhomogeneous hard-sphere fluid: Application to interfacial adsorption. *Phys. Rev. A* **42**, 3382–3387 (1990).
10. E. Kierlik, M. L. Rosinberg, Density-functional theory for inhomogeneous fluids: Adsorption of binary mixtures. *Phys. Rev. A* **44**, 5025–5037 (1991).
11. S. Phan, E. Kierlik, M. L. Rosinberg, B. Bildstein, G. Kahl, Equivalence of two free-energy models for the inhomogeneous hard-sphere fluid. *Phys. Rev. E* **48**, 618–620 (1993).
12. R. Roth, R. Evans, A. Lang, G. Kahl, Fundamental measure theory for hard-sphere mixtures revisited: The White Bear version. *J. Phys.: Condens. Matter* **14**, 12063–12078 (2002).
13. H. Hansen-Goos, R. Roth, Density functional theory for hard-sphere mixtures: The White Bear version mark II. *J. Phys.: Condens. Matter* **18**, 8413–8425 (2006).
14. R. Roth, Fundamental measure theory for hard-sphere mixtures: A review. *J. Phys.: Condens. Matter* **22**, 063102 (2010).
15. P. Tarazona, Density functional for hard sphere crystals: A fundamental measure approach. *Phys. Rev. Lett.* **84**, 694–697 (2000).
16. P. Tarazona, J. Cuesta, Y. Martínez-Ratón, “Density functional theories of hard particle systems” in *Theory and Simulation of Hard-Sphere Fluids and Related Systems* (Springer, Berlin, 2008), pp. 247–341.
17. J. K. Percus, G. J. Yevick, Analysis of classical statistical mechanics by means of collective coordinates. *Phys. Rev.* **110**, 1–13 (1958).
18. Y. Rosenfeld, Scaled field particle theory of the structure and the thermodynamics of isotropic hard particle fluids. *J. Chem. Phys.* **89**, 4272–4287 (1988).
19. S. R. Nagel, Experimental soft-matter science. *Rev. Mod. Phys.* **89**, 025002 (2017).
20. R. Evans, D. Frenkel, M. Dijkstra, From simple liquids to colloids and soft matter. *Phys. Today* **72**, 38–39 (2019).
21. J. L. Barrat, J. P. Hansen, *Basic Concepts for Simple and Complex Liquids* (Cambridge University Press, 2003).
22. V. Molinero, E. B. Moore, Water modeled as an intermediate element between carbon and silicon. *J. Phys. Chem. B* **113**, 4008–4016 (2008).
23. M. K. Coe, R. Evans, N. B. Wilding, The coexistence curve and surface tension of a monatomic water model. *J. Chem. Phys.* **156**, 154505 (2022).
24. M. K. Coe, R. Evans, N. B. Wilding, Density depletion and enhanced fluctuations in water near hydrophobic solutes: Identifying the underlying physics. *Phys. Rev. Lett.* **128**, 045501 (2022).
25. M. K. Coe, R. Evans, N. B. Wilding, Understanding the physics of hydrophobic solvation. *J. Chem. Phys.* **158**, 034508 (2023).
26. D. Martín-Jiménez, E. Chacon, P. Tarazona, R. García, Atomically resolved three-dimensional structures of electrolyte aqueous solutions near a solid surface. *Nat. Commun.* **7**, 12164 (2016).
27. J. Hernández-Muñoz, E. Chacón, P. Tarazona, Density functional analysis of atomic force microscopy in a dense fluid. *J. Chem. Phys.* **151**, 034701 (2019).
28. P. Cats, R. Evans, A. Härtel, R. van Roij, Primitive model electrolytes in the near and far field: Decay lengths from DFT and simulations. *J. Chem. Phys.* **154**, 124504 (2021).
29. P. S. Clegg, Characterising soft matter using machine learning. *Soft Matter* **17**, 3991–4005 (2021).
30. M. Dijkstra, E. Luijten, From predictive modelling to machine learning and reverse engineering of colloidal self-assembly. *Nat. Mater.* **20**, 762–773 (2021).
31. E. Boattini, M. Dijkstra, L. Filion, Unsupervised learning for local structure detection in colloidal systems. *J. Chem. Phys.* **151**, 154901 (2019).
32. G. Campos-Villalobos, E. Boattini, L. Filion, M. Dijkstra, Machine learning many-body potentials for colloidal systems. *J. Chem. Phys.* **155**, 174902 (2021).
33. G. Campos-Villalobos, G. Giunta, S. Marín-Aguilar, M. Dijkstra, Machine-learning effective many-body potentials for anisotropic particles using orientation-dependent symmetry functions. *J. Chem. Phys.* **157**, 024902 (2022).
34. T. Santos-Silva, P. I. C. Teixeira, C. Anquetil-Deck, D. J. Cleaver, Neural-network approach to modeling liquid crystals in complex confinement. *Phys. Rev. E* **89**, 053316 (2014).
35. S. C. Lin, M. Oettel, A classical density functional from machine learning and a convolutional neural network. *SciPost Phys.* **6**, 025 (2019).
36. S. C. Lin, G. Martius, M. Oettel, Analytical classical density functionals from an equation learning network. *J. Chem. Phys.* **152**, 021102 (2020).
37. P. Yatsyshin, S. Kalliadas, A. B. Duncan, Physics-constrained Bayesian inference of state functions in classical density-functional theory. *J. Chem. Phys.* **156**, 074105 (2022).
38. P. Cats *et al.*, Machine-learning free-energy functionals using density profiles from simulations. *APL Mater.* **9**, 031109 (2021).
39. D. de las Heras, T. Zimmermann, F. Sannüller, S. Hermann, M. Schmidt, Perspective: How to overcome dynamical density functional theory. *J. Phys.: Condens. Matter* **35**, 271501 (2023).
40. M. Schmidt, J. M. Brader, Power functional theory for Brownian dynamics. *J. Chem. Phys.* **138**, 214101 (2013).
41. M. Schmidt, Power functional theory for many-body dynamics. *Rev. Mod. Phys.* **94**, 015007 (2022).
42. S. Hermann, M. Schmidt, Noether’s theorem in statistical mechanics. *Commun. Phys.* **4**, 176 (2021).
43. F. Chollet, *Deep Learning with Python* (Manning Publications, 2017).
44. R. D. Groot, J. P. van der Eerden, N. M. Faber, The direct correlation function in hard sphere fluids. *J. Chem. Phys.* **87**, 2263–2270 (1987).
45. S. R. Dubey, S. K. Singh, B. B. Chaudhuri, Activation functions in deep learning: A comprehensive survey and benchmark. *Neurocomputing* **503**, 92–108 (2022).
46. X. Glorot, A. Bordes, Y. Bengio, “Deep sparse rectifier neural networks” in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (JMLR Workshop and Conference Proceedings)* (2011), pp. 315–323.
47. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization (2014).
48. D. Henderson, E. W. Grundke, Direct correlation function: Hard sphere fluid. *J. Chem. Phys.* **63**, 601–607 (1975).
49. J. M. Brader, M. Schmidt, Free power dissipation from functional line integration. *Mol. Phys.* **113**, 2873–2880 (2015).
50. J. Kolafa, S. Labík, A. Malijevský, Accurate equation of state of the hard sphere fluid in stable and metastable regions. *Phys. Chem. Chem. Phys.* **6**, 2335–2340 (2004).
51. N. F. Carnahan, K. E. Starling, Equation of state for nonattracting rigid spheres. *J. Chem. Phys.* **51**, 635–636 (1969).
52. F. Höfling, S. Dietrich, Finite-size corrections for the static structure factor of a liquid slab with open boundaries. *J. Chem. Phys.* **153**, 054119 (2020).
53. Y. Rosenfeld, M. Schmidt, H. Löwen, P. Tarazona, Fundamental-measure free-energy density functional for hard spheres: Dimensional crossover and freezing. *Phys. Rev. E* **55**, 4245–4263 (1997).
54. D. Frenkel, B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Elsevier Science & Technology, 2023).
55. R. L. Davidchack, B. B. Laird, R. Roth, Hard spheres at a planar hard wall: Simulations and density functional theory. *Condens. Matter Phys.* **19**, 23001 (2016).
56. F. Sannüller, S. Hermann, M. Schmidt, Comparative study of force-based classical density functional theory. *Phys. Rev. E* **107**, 034109 (2023).
57. D. de las Heras, M. Schmidt, The phase stacking diagram of colloidal mixtures under gravity. *Soft Matter* **9**, 8636 (2013).
58. L. D. Site *et al.*, Molecular dynamics of open systems: Construction of a mean-field particle reservoir. *Adv. Theory Simul.* **2**, 1900014 (2019).
59. L. A. Baptista *et al.*, Density-functional-theory approach to the Hamiltonian adaptive resolution simulation method. *J. Phys.: Condens. Matter* **33**, 184003 (2021).
60. F. Schmid, Editorial, Multiscale simulation methods for soft matter systems. *J. Phys.: Condens. Matter* **34**, 160401 (2022).
61. G. E. Karniadakis *et al.*, Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440 (2021).
62. A. J. Archer, B. Chacko, R. Evans, The standard mean-field treatment of inter-particle attraction in classical DFT is better than one might expect. *J. Chem. Phys.* **147**, 034501 (2017).
63. S. M. Tschopp, H. D. Vuijk, A. Sharma, J. M. Brader, Mean-field theory of inhomogeneous fluids. *Phys. Rev. E* **102**, 042140 (2020).
64. M. Schmidt, Density-functional theory for soft interactions by dimensional crossover. *Phys. Rev. E* **60**, R6291–R6294 (1999).
65. M. Schmidt, Fluid structure from density-functional theory. *Phys. Rev. E* **62**, 4976–4981 (2000).
66. K. L. Finster, E. J. Krebs, C. J. May, P. A. Kreitzberg, D. Roundy, Soft fundamental measure theory functional for the Weeks-Chandler-Andersen repulsive potential. *Phys. Rev. E* **106**, 064134 (2022).
67. G. Jeanmairet, M. Levesque, R. Vuilleumier, D. Borgis, Molecular density functional theory of water. *J. Phys. Chem.* **4**, 619–624 (2013).
68. L. Ding, M. Levesque, D. Borgis, L. Belloni, Efficient molecular density functional theory using generalized spherical harmonics expansions. *J. Chem. Phys.* **147**, 094107 (2017).
69. T. S. Cohen, M. Welling, Group equivariant convolutional networks (2016).
70. M. Weiler, M. Geiger, M. Welling, W. Boomsma, T. Cohen, 3D steerable CNNs: Learning rotationally equivariant features in volumetric data (2018).
71. M. Finzi, S. Stanton, P. Izmailov, A. G. Wilson, Generalizing convolutional neural networks for equivariance to Lie groups on arbitrary continuous data (2020).
72. V. G. Satorras, E. Hoogeboom, M. Welling, “E(n) equivariant graph neural networks” in *Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research*, M. Meila, T. Zhang, Eds. (PMLR, 2021), vol. 139, pp. 9323–9332.
73. S. Batzner *et al.*, E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **13**, 2453 (2022).
74. S. Batzner, A. Musaelian, B. Kozinsky, Advancing molecular simulation with equivariant interatomic potentials. *Nat. Rev. Phys.* **5**, 437–438 (2023).
75. A. Musaelian *et al.*, Learning local equivariant representations for large-scale atomistic dynamics. *Nat. Commun.* **14**, 579 (2023).
76. R. Pederson, B. Kalita, K. Burke, Machine learning and density functional theory. *Nat. Rev. Phys.* **4**, 357–358 (2022).
77. Y. Zhou, J. Wu, S. Chen, G. Chen, Toward the exact exchange-correlation potential: A three-dimensional convolutional neural network construct. *J. Phys. Chem. Lett.* **10**, 7264–7269 (2019).
78. R. Nagai, R. Akashi, S. Sasaki, S. Tsuneyuki, Neural-network Kohn-Sham exchange-correlation potential and its out-of-training transferability. *J. Chem. Phys.* **148**, 241737 (2018).
79. J. Schmidt, C. L. Benavides-Riveros, M. A. L. Marques, Machine learning the physical nonlocal exchange-correlation functional of density-functional theory. *J. Phys. Chem. Lett.* **10**, 6425–6431 (2019).
80. R. Nagai, R. Akashi, O. Sugino, Completing density functional theory by machine learning hidden messages from molecules. *npj Comput. Mater.* **6**, 43 (2020).
81. L. Li *et al.*, Kohn-Sham equations as regularizer: Building prior knowledge into machine-learned physics. *Phys. Rev. Lett.* **126**, 036401 (2021).
82. H. Li *et al.*, Deep-learning density functional theory Hamiltonian for efficient ab initio electronic-structure calculation. *Nat. Comput. Sci.* **2**, 367–377 (2022).
83. J. Wang, Y. Wang, R. X. Xu, G. Chen, X. Zheng, A semilocal machine-learning correction to density functional approximations. *J. Chem. Phys.* **158**, 154107 (2023).
84. F. Sannüller, S. Hermann, D. de las Heras, M. Schmidt, Neural functional theory for inhomogeneous fluids: Fundamentals and applications. Zenodo. <https://zenodo.org/records/8380003>. Deposited 26 September 2023.