

Determining the Chemical Potential via Universal Density Functional Learning

Florian Sammüller^{✉*} and Matthias Schmidt^{✉†}*Theoretische Physik II, Physikalisches Institut, Universität Bayreuth, D-95447 Bayreuth, Germany*

(Received 18 June 2025; accepted 23 December 2025; published 10 February 2026)

We demonstrate that the machine learning of density functionals allows one to determine simultaneously the equilibrium chemical potential across simulation datasets of inhomogeneous classical fluids. Minimization of a loss function based on an Euler-Lagrange equation yields both the universal one-body direct correlation functional, which is represented locally by a neural network, as well as the system-specific unknown chemical potential values. The method can serve as an efficient alternative to conventional computational techniques of measuring the chemical potential. It also facilitates using canonical data from Brownian dynamics, molecular dynamics, or Monte Carlo simulations as a basis for constructing neural density functionals, which are fit for accurate multiscale prediction of soft matter systems in equilibrium.

DOI: [10.1103/7bqn-y2d7](https://doi.org/10.1103/7bqn-y2d7)

Various different approaches to understanding the chemical potential μ have been put forward [1–3]. From a fundamental statistical mechanical perspective, μ acts as an additive constant to the potential energy of each particle [4]. The corresponding statistical weight for each microstate of N particles is $e^{-\beta(H-\mu N)}/\Xi$, where H denotes the Hamiltonian, Ξ is the grand partition sum, and $\beta = 1/(k_B T)$ with absolute temperature T and the Boltzmann constant k_B . In computer simulations, Widom's test particle insertion method [5,6] and its powerful variants and generalizations (see, e.g., Refs. [7–13]) allow one to express the chemical potential as a canonical average, which circumvents the need to represent the grand ensemble directly [14]. Specifically, the nontrivial excess (over ideal gas) chemical potential is obtained as $\beta\mu_{\text{exc}} = -\ln\langle\exp(-\beta\epsilon)\rangle$, where the angle brackets denote a canonical thermal equilibrium average with N particles and ϵ is the energy change caused by adding an $(N + 1)$ th particle. Generating values for ϵ is readily performed in simulations via virtual particle insertions.

Testing and developing methods that are applicable to hard sphere systems [15,16] is important due to the particular features arising from purely entropic behavior [17–20]. Via analyzing colloidal particle coordinates obtained from optical microscopy, Dullens, Aarts, and Kegel [21] determined the chemical potential in their

experimental hard sphere system. Computational particle insertion can be performed in a variety of ways, including via growth from an added fourth dimension [22]. Widom insertion also generalizes to more complex systems, such as polymers [23]. Furthermore, the chemical potential was shown to be accessible via adaptive resolution [24–27] and parallel computing techniques [28,29]. Its role in relevant questions of phase coexistence [30] and of finite size effects [31] was addressed. While much of the above work concerns bulk fluids, several insightful studies were specifically targeted at spatially inhomogeneous systems [32–36], which are pertinent to tackle the wide range of adsorption and interfacial phenomena in soft matter.

Machine learning is currently applied across a rapidly growing scope of problems in statistical mechanics, ranging from the characterization of soft matter [37], reverse engineering of self-assembly [38], and local structure detection [39] to the investigation of many-body potentials [40,41]. Notably, a surge of machine learning techniques has recently been developed in the context of classical density functional theory [42–70]. Besides the prime goal of approximating the underlying density functional mappings, various methods have also proven feasible for addressing complex observables [58,59], inverse design problems [52,53,60], and the dynamics in nonequilibrium systems [52,53].

We recall that classical density functional theory is founded on a formally exact variational principle of the grand potential [71], which is expressed as a functional of the density profile $\rho(\mathbf{r})$, where \mathbf{r} denotes (generic) position. Analogously to its quantum mechanical counterpart [72,73], this theoretical underpinning facilitates the development of highly accurate and efficient machine learning techniques. Specifically, the rigorous functional relationships lend themselves very naturally to be represented by

*Contact author: florian.sammueler@uni-bayreuth.de†Contact author: matthias.schmidt@uni-bayreuth.de

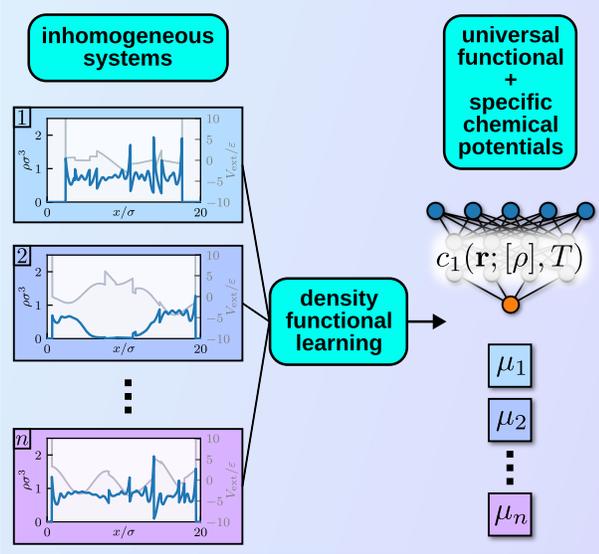


FIG. 1. Work flow for determining chemical potential values across a dataset of n inhomogeneous systems via density functional learning. Data from one simulation run k consist of the temperature value T_k , the external potential $V_{\text{ext}}^{(k)}(\mathbf{r})$ (gray lines), and the density profile $\rho_k(\mathbf{r})$ (blue lines)—in one spatial dimension or for planar symmetry $\mathbf{r} = x$. Supervised training based on the loss function (3) determines both the universal one-body direct correlation functional $c_1(\mathbf{r}; [\rho], T)$, which is represented locally as a neural network (see Appendix A), as well as the system-specific chemical potential values μ_k .

neural networks, and they enable a tight integration of modern computational tools, in particular, automatic differentiation, for analyzing a wide range of physical behavior [74].

All currently available methods for learning classical density functionals rely on the explicit access to grand canonical simulation data, which entails known (prescribed) chemical potential values for each system in a given dataset of inhomogeneous fluid equilibria. In this Letter, we present an augmentation of density functional learning, which does not require the chemical potential values to be known. We show that hitherto unknown chemical potentials can instead be determined simultaneously for all systems during the training of neural density functionals without any additional computational effort and crucially without resorting back to many-body simulation techniques. The benefit of our approach is twofold: (i) The method constitutes a high-throughput scheme for measuring chemical potentials across large datasets of inhomogeneous fluids. (ii) Constructing neural density functionals is no longer restricted to using simulation data with full grand canonical information. Both of these aspects are relevant in practice, and we describe possible applications below. Making density functional learning more versatile is particularly pertinent in light of the arguably much more widespread use of simulation techniques in the NVT and NpT ensembles compared to grand canonical schemes.

Concretely, our method facilitates the direct utilization of data from molecular or Brownian dynamics simulations for training neural functionals, provided that ensemble differences in the acquired density profiles remain negligible. An illustration of the central work flow is shown in Fig. 1.

We first give an overview of fundamentals. A central object of classical density functional theory is the one-body direct correlation functional $c_1(\mathbf{r}; [\rho], T)$. Besides its functional dependence on the density profile $\rho(\mathbf{r})$, denoted by brackets, the temperature T arises as an additional parametric dependence for general (beyond hard-core) fluid types. We keep the following presentation general but drop the dependence on T if athermal systems or applications for fixed temperature are considered. For a given type of fluid, as specified by its intrinsic Hamiltonian, $c_1(\mathbf{r}; [\rho], T)$ is a unique functional, independent of the external potential [71]. It is, hence, universally applicable for predicting equilibrium states of the considered fluid in arbitrary inhomogeneous environments [75]. Assuming that the universal functional $c_1(\mathbf{r}; [\rho], T)$ is known, self-consistently solving a corresponding Euler-Lagrange equation,

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

yields a prediction for the density profile $\rho(\mathbf{r})$ that is generated by the external potential $V_{\text{ext}}(\mathbf{r})$ at the given state point. Obtaining a numerical solution of Eq. (1) is typically both straightforward and efficient, even for very large systems with high spatial resolution [54]. Because of the fundamental role of $c_1(\mathbf{r}; [\rho], T)$, recent machine learning approaches have focused on representing this functional in terms of a neural network; an overview of feasible techniques for the construction of neural density functionals is given in Appendix A.

Equation (1) is also key for the analysis of a dataset of inhomogeneous fluid states, and it will, hence, serve as the basis for training neural density functionals. As a template, we first consider grand canonical simulations, where the thermodynamic state point is characterized by known system-specific values of T and μ . Upon imposing a spatially inhomogeneous external potential $V_{\text{ext}}(\mathbf{r})$, the resulting density profile $\rho(\mathbf{r})$ becomes inhomogeneous, in general. Its microscopic definition is $\rho(\mathbf{r}) = \langle \sum_i \delta(\mathbf{r} - \mathbf{r}_i) \rangle$, where the angle brackets indicate an equilibrium grand ensemble average, the sum runs over all particles $i = 1, \dots, N$ with positions \mathbf{r}_i , and $\delta(\cdot)$ denotes the Dirac distribution. This average can be measured in simulations via standard histogram-counting methods or with more advanced reduced-variance (force-sampling) algorithms [76–78].

We henceforth consider n individual simulation runs which contribute to the dataset and assign to each result the label $k = 1, \dots, n$ to keep track of the respective simulation origin. Rearranging Eq. (1) yields

$$c_1^{(k)}(\mathbf{r}) = \ln \rho_k(\mathbf{r}) + \beta_k V_{\text{ext}}^{(k)}(\mathbf{r}) - \beta_k \mu_k, \quad (2)$$

which makes all quantities on the right-hand side known. Via Eq. (2), the one-body direct correlation function $c_1^{(k)}(\mathbf{r})$ that corresponds to the grand canonical simulation run k is readily computed in a postprocessing step on the acquired simulation data for each position \mathbf{r} where $\rho_k(\mathbf{r}) > 0$. The relationship $\rho_k(\mathbf{r}), T_k \rightarrow c_1^{(k)}(\mathbf{r})$ can, hence, be constructed explicitly for all simulated systems if all values of μ_k are known. As classical density functional theory ascertains the universality of this functional mapping, training a neural network on the thus-acquired dataset is feasible in order to obtain a highly accurate numerical representation of $c_1(\mathbf{r}; [\rho], T)$. Specifically, we use the local one-body correlation learning scheme [54–56,61] and provide further details in Appendix A.

We now modify the density functional learning in order to not require the chemical potential values to be known. Hence, our aim is to work solely with input data $T_k, V_{\text{ext}}^{(k)}(\mathbf{r})$, and $\rho_k(\mathbf{r})$, as are directly available already in canonical simulations. Our treatment of the set of the then-unknown μ_k is similar to treating the parameters of the neural network: Both are to be determined during the training procedure. Additionally to the neural network which represents $c_1(\mathbf{r}; [\rho], T)$, we therefore prepare an array of n numerical variables which correspond individually to the chemical potentials of the n simulations. The neural network parameters and the chemical potential values are simultaneously optimized via standard backpropagation, as illustrated in Fig. 1 and as described in the following.

The method rests on the fact that $\beta\mu$ merely constitutes an additive offset to $c_1(\mathbf{r}; [\rho], T)$ in Eq. (1), which is identical for all positions \mathbf{r} in a given system. Thus, Eq. (1) serves as a simultaneous matching condition for both $c_1(\mathbf{r}; [\rho_k], T_k)$ and μ_k for each simulation run k in the dataset. Therefore, we, define the loss function

$$\mathcal{L}_{\text{EL}} = \sum_{k=1}^n \left\| \ln \rho_k(\mathbf{r}) + \beta_k V_{\text{ext}}^{(k)}(\mathbf{r}) - \beta_k \mu_k - c_1(\mathbf{r}; [\rho_k], T_k) \right\|_2^2, \quad (3)$$

which measures via the spatial norm $\|\cdot\|_2$ how strongly the Euler-Lagrange equation (1) is violated across all members $k = 1, \dots, n$ of the dataset. Both the functional form of $c_1(\mathbf{r}; [\rho], T)$ as well as the values μ_k thereby enter as predictions and, hence, as the targets to be optimized. In practice, all profiles are provided numerically at discretized positions \mathbf{r} . The implementation of Eq. (3) is straightforward [79], as its form allows utilizing the standard mean squared error loss function readily available in common machine learning libraries.

Clearly, $\mathcal{L}_{\text{EL}} \geq 0$, with equality holding only if Eq. (1) is satisfied across the entire dataset. Hence, minimization of the loss function (3) determines simultaneously the values

μ_k and the functional $c_1(\mathbf{r}; [\rho], T)$, where the latter is represented by the neural network parameters. We emphasize two central points which rationalize the successful convergence (see below) of this minimization: (i) The functional mapping $c_1(\mathbf{r}; [\rho], T)$ is universal for the considered fluid type [71]; i.e., it is independent of chemical and external potential. (ii) For each simulation run k , while the functional relationship contributes individually at every (discretized) position \mathbf{r} throughout the system to the loss function (3), the chemical potential necessarily corresponds to the same system-specific value μ_k . Thus, the values μ_k can be recovered consistently across the whole dataset.

As a practical detail, the minimization of Eq. (3) determines the resulting functional $c_1(\mathbf{r}; [\rho], T)$ and the set of values $\beta_k \mu_k$ only up to a global additive constant \tilde{c}_1 . This follows trivially from substituting $\beta_k \mu_k \rightarrow \beta_k \mu_k + \tilde{c}_1$ and $c_1(\mathbf{r}; [\rho_k], T_k) \rightarrow c_1(\mathbf{r}; [\rho_k], T_k) - \tilde{c}_1$ in Eq. (3), which leaves \mathcal{L}_{EL} unchanged. This global offset can be corrected *a posteriori* by demanding that $c_1(\mathbf{r}; [\rho = 0], T) = 0$. In applications, the neural functional must be evaluated for vanishing density input, and this offset needs to be subtracted from the inferred value at target density $\rho(\mathbf{r})$ to obtain the correct result for $c_1(\mathbf{r}; [\rho], T)$. Similarly, the recovered values μ_k have to be corrected by the same (negative) offset.

For an initial proof of principle, we turn to (athermal) one-dimensional hard rods with particle size σ , where neural density functionals deliver highly accurate results [55,56], and with Percus' exact analytical free energy functional [80] serving as the ultimate benchmark. While we use grand canonical Monte Carlo simulations, which comprise the full dataset $\mu_k, T_k, \rho_k(x)$, and $V_{\text{ext}}^{(k)}(x)$, where $k = 1, \dots, 512$ and x denotes one-dimensional position, we deliberately disregard all values μ_k during training. We proceed instead via the above methodology using the Euler-Lagrange-based loss function (3) to determine these values along with the functional $c_1(x; [\rho])$. Our goal is, hence, to recover via the training all values μ_k , which can be verified *a posteriori* against the (known) grand canonical Monte Carlo input parameters. We observe that the values μ_k converge already after few training epochs. Comparison with the available reference data for μ_k , upon correcting the global offset \tilde{c}_1 , confirms excellent agreement, as shown in Fig. 2(a), which corroborates that unknown chemical potentials can be recovered via the present procedure. Moreover, utilizing the resulting neural functional $c_1(x; [\rho])$ in predictions yields density profiles which closely match those obtained from the exact Percus functional, as verified for the case of hard-wall confinement.

In order to demonstrate the prowess of the method, we next consider the three-dimensional truncated Lennard-Jones fluid, which is specified by the pairwise interparticle potential $\phi(r) = 4\epsilon[(\sigma/r)^{12} - (\sigma/r)^6]$ for particle distance $r < r_c = 2.5\sigma$ and $\phi(r) = 0$ otherwise, where ϵ and σ set the energy and length scale, respectively. The grand

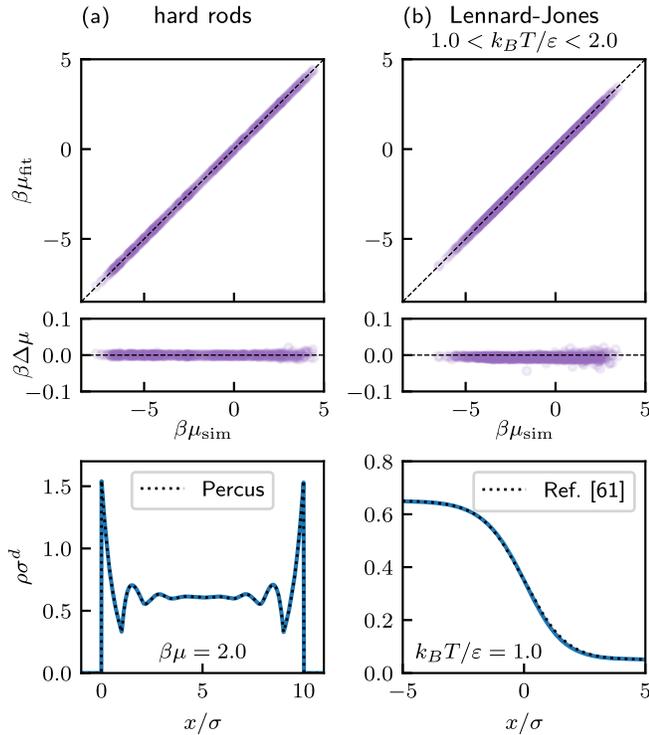


FIG. 2. Predictions for the scaled chemical potential values $\beta\mu_{\text{fit}}$ (top row) for datasets of hard rods in $d = 1$ dimension (a) and the truncated Lennard-Jones fluid in $d = 3$ dimensions and planar symmetry (b). The Lennard-Jones fluid is addressed via thermal training [61] in the temperature range $1.0 < k_B T/\varepsilon < 2.0$. The results show excellent agreement with the simulation reference $\beta\mu_{\text{sim}}$; see the consistently small difference $\beta\Delta\mu = \beta\mu_{\text{fit}} - \beta\mu_{\text{sim}}$ (middle row). The trained neural correlation functionals $c_1(x; [\rho])$ (a) and $c_1(x; [\rho], T)$ (b) are directly applicable for the calculation of density profiles (bottom row, solid blue lines). We consider hard-wall confinement in a planar slit of length 10σ for hard rods (a) and compare the prediction to Percus' exact density functional result [80] (dotted black line). For the thermally trained functional of the Lennard-Jones fluid (b), we predict phase coexistence and show an exemplary density profile of the liquid-gas interface, which is compared to the neural result of Ref. [61] (dotted black line).

canonical reference dataset consists of $n = 882$ simulation results at varying temperatures $1.0 < k_B T/\varepsilon < 2.0$ such that the neural functional $c_1(x; [\rho], T)$ acquires additional parametric dependence on temperature [61]. Inhomogeneous external potentials are imposed in planar geometry with x denoting the spatial coordinate along the inhomogeneous direction. As before, all values μ_k are, though available, crucially not utilized in the training. Instead, the chemical potentials are recovered during density functional learning, and the resulting values closely match the reference also in this case, as shown in Fig. 2(b). Additionally, the thermally trained neural functional is fit for application, which we demonstrate for the prototypical problem of determining the interfacial density profile of coexisting liquid and gas phases [61]. In particular, training on the basis of Eq. (3) without chemical potential values causes no noticeable decrease in

the accuracy of the neural functional as compared to direct training via Eq. (2) with known values of μ_k . We recall that both methods deliver neural functionals that remain robust for predicting fluid equilibria not featured explicitly in the training dataset, such as the shown case of liquid-gas coexistence [61]. Further results for the hard sphere and isothermal Lennard-Jones fluids are presented in Appendix B. Code and data to reproduce the findings of this work are openly accessible [79].

We next turn to canonical training and consider the truncated Lennard-Jones fluid at fixed temperature $k_B T = 1.5\varepsilon$ [54,57]. The dataset comprises $n = 400$ individual canonical Monte Carlo simulation runs with randomized external potential landscapes and randomized particle numbers in the range $100 < N < 1500$ yielding mean densities $0.05 < \bar{\rho}\sigma^3 < 0.75$. Because of the sufficiently large system size $10 \times 10 \times 20\sigma^3$ and large particle numbers, we expect ensemble differences in the density profiles to be negligible and, hence, to obtain a suitable dataset for determining a universal functional mapping $c_1(x; [\rho])$. The training proceeds on the available data for $V_{\text{ext}}^{(k)}(x)$ and the canonical density profiles $\rho_k(x)$, where $k = 1, \dots, n$, and we observe that the minimization of the loss function (3) converges. As no reference for chemical potential values is available, we perform an indirect verification by testing the predictions of the neural density functional $c_1(x; [\rho])$, reasoning that if this functional mapping is represented correctly, then the chemical potential values must follow accordingly from the successful minimization of Eq. (3). Figure 3 shows a self-consistent density profile for planar hard-wall confinement obtained via the canonically trained neural functional. For comparison, we also show the prediction of a neural functional trained on grand canonical data (see Appendix B) as well as a density profile from direct grand canonical Monte Carlo simulation. All results agree with each other, thus confirming the feasibility of constructing neural density functionals based on canonical simulation data.

In summary, the presented method constitutes an efficient means of determining chemical potential values for a given dataset of n inhomogeneous fluid states. The sole input quantities are the temperatures T_k , the external potentials $V_{\text{ext}}^{(k)}(\mathbf{r})$, and the corresponding density profiles $\rho_k(\mathbf{r})$ of each system $k = 1, \dots, n$, which are directly accessible already in canonical simulation methods such as molecular or Brownian dynamics. The viability of our framework fundamentally rests on the existence and uniqueness of the classical density functional mapping $T, \rho(\mathbf{r}) \rightarrow V_{\text{ext}}(\mathbf{r})$ [71]. In principle, a dataset of these quantities for various inhomogeneities should, therefore, be sufficient for recovering the many-body physics of the fluid model under consideration, implying that unknown chemical potentials can formally be determined *a posteriori*. To put this task into practice, we have exploited the universal nature of the one-body direct correlation functional

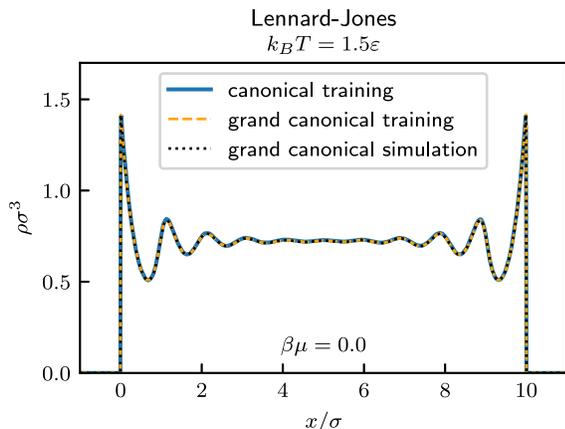


FIG. 3. Results for the density profile (solid blue line) of the truncated Lennard-Jones fluid in planar hard-wall confinement obtained from training a neural functional with data from canonical Monte Carlo simulations with unknown chemical potential values. For comparison, we have also performed training with grand canonical data (dashed orange line; see also Appendix B and Fig. 4) and take as ultimate reference the density profile from a grand canonical Monte Carlo simulation (dotted black line). The agreement of all three routes confirms the accuracy of the prediction and the feasibility of neural density functional construction based on canonical data.

$c_1(\mathbf{r}; [\rho], T)$. Density functional learning of a neural-network-based representation of $c_1(\mathbf{r}; [\rho], T)$ via the loss function (3) then yields the missing system-specific chemical potentials μ_k essentially as a mere additive offset. We reiterate that not a single value of μ_k needs to be known *a priori* and that our first-principles-based method is conceptually different from more common applications of machine learning, where the underlying dataset usually features the target quantity directly.

The application of our method is particularly beneficial for large datasets, which arise, for instance, in solvation studies that are relevant for various biochemical processes and for computer-aided drug discovery [81–84]. Direct means of measuring the chemical potential, e.g., using Widom’s insertion method, incur additional computational cost that scales linearly with the number n of considered systems, which might be prohibitive for such data-intensive tasks. In contrast, our method rests only on the availability of density profiles, from which the chemical potentials as well as further results can be obtained without resorting back to many-body simulation techniques. As merely a machine-learning-based postprocessing step is required, there is no relevant computational overhead besides the comparatively inexpensive density functional learning via Eq. (3). Recall further that the trained neural functional can readily be utilized for predictions of various thermodynamic and structural properties [54,61], such as (solvation) free energies.

We conclude with several practical aspects and conceptual points. Because of the underlying mechanism of

density functional learning, the functional mapping needs to be probed sufficiently by the provided dataset such that training of a neural functional succeeds. Datasets with only very few systems or with weak inhomogeneities may hamper this machine learning task; the reliability of predictions for reduced dataset sizes is investigated in Appendix C. Pair-correlation regularization [57,68] could remedy the limitations of insufficient inhomogeneous data to a certain extent by additionally matching the neural functional to two-body correlation functions in the bulk fluid. Lastly, we emphasize that the temperature is required as a known quantity in the dataset. It would be interesting in future work to evaluate the feasibility of using purely microcanonical data for training neural functionals and for recovering temperature values [85] but also to investigate ensemble differences and canonical density functionals [86] as well as connections to metadensity functional theory [60].

Acknowledgments—We thank Robert Evans for valuable discussions. Some of the calculations were performed using the emil and festus clusters of the Bayreuth Center for High Performance Computing funded by the DFG (Deutsche Forschungsgemeinschaft) under Projects No. 422127126 and No. 523317330. This work is supported by the DFG under Project No. 551294732.

Data availability—The data that support the findings of this article are openly available [79].

- [1] G. Cook and R. H. Dickerson, Understanding the chemical potential, *Am. J. Phys.* **63**, 737 (1995).
- [2] R. Baierlein, The elusive chemical potential, *Am. J. Phys.* **69**, 423 (2001).
- [3] G. Job and F. Herrmann, Chemical potential—A quantity in search of recognition, *Eur. J. Phys.* **27**, 353 (2006).
- [4] J.-P. Hansen and I. R. McDonald, *Theory of Simple Liquids: With Applications to Soft Matter*, 4th ed. (Academic Press, Amsterdam, 2013).
- [5] B. Widom, Some topics in the theory of fluids, *J. Chem. Phys.* **39**, 2808 (1963).
- [6] J. L. Jackson and L. S. Klein, Potential distribution method in equilibrium statistical mechanics, *Phys. Fluids* **7**, 228 (1964).
- [7] I. Nezbeda and J. Kolafa, A new version of the insertion particle method for determining the chemical potential by Monte Carlo simulation, *Mol. Simul.* **5**, 391 (1991).
- [8] J. Powles, W. Evans, and N. Quirke, Non-destructive molecular-dynamics simulation of the chemical potential of a fluid, *Mol. Phys.* **46**, 1347 (1982).
- [9] K. Shing and K. Gubbins, The chemical potential from computer simulation: Test particle method with umbrella sampling, *Mol. Phys.* **43**, 717 (1981).
- [10] K. Shing and K. Gubbins, The chemical potential in non-ideal liquid mixtures: Computer simulation and theory, *Mol. Phys.* **49**, 1121 (1983).

- [11] B. Smit and D. Frenkel, Calculation of the chemical potential in the Gibbs ensemble, *Mol. Phys.* **68**, 951 (1989).
- [12] D. A. Kofke and P. T. Cummings, Quantitative comparison and optimization of methods for evaluating the chemical potential by molecular simulation, *Mol. Phys.* **92**, 973 (1997).
- [13] N. Lu, J. K. Singh, and D. A. Kofke, Appropriate methods to combine forward and reverse free-energy perturbation averages, *J. Chem. Phys.* **118**, 2977 (2003).
- [14] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications*, 3rd ed. (Academic Press, Amsterdam, 2023).
- [15] A. Santos, S. B. Yuste, and M. López De Haro, Structural and thermodynamic properties of hard-sphere fluids, *J. Chem. Phys.* **153**, 120901 (2020).
- [16] C. P. Royall, P. Charbonneau, M. Dijkstra, J. Russo, F. Smallenburg, T. Speck, and C. Valeriani, Colloidal hard spheres: Triumphs, challenges, and mysteries, *Rev. Mod. Phys.* **96**, 045003 (2024).
- [17] P. Attard, Simulation of the chemical potential and the cavity free energy of dense hard-sphere fluids, *J. Chem. Phys.* **98**, 2225 (1993).
- [18] D. M. Heyes and A. Santos, Chemical potential of a test hard sphere of variable size in a hard-sphere fluid, *J. Chem. Phys.* **145**, 214504 (2016).
- [19] D. M. Heyes and A. Santos, Chemical potential of a test hard sphere of variable size in hard-sphere fluid mixtures, *J. Chem. Phys.* **148**, 214503 (2018).
- [20] R. L. Davidchack and B. B. Laird, Chemical potential and surface free energy of a hard spherical particle in hard-sphere fluid over the full range of particle diameters, *J. Chem. Phys.* **157**, 074701 (2022).
- [21] R. P. A. Dullens, D. G. A. L. Aarts, and W. K. Kegel, Direct measurement of the free energy by optical microscopy, *Proc. Natl. Acad. Sci. U.S.A.* **103**, 529 (2006).
- [22] L. Belloni, Non-equilibrium hybrid insertion/extraction through the 4th dimension in grand-canonical simulation, *J. Chem. Phys.* **151**, 021101 (2019).
- [23] N. B. Wilding and M. Müller, Accurate measurements of the chemical potential of polymeric systems by Monte Carlo simulation, *J. Chem. Phys.* **101**, 4324 (1994).
- [24] H. Wang, C. Hartmann, C. Schütte, and L. Delle Site, Grand-canonical-like molecular-dynamics simulations by using an adaptive-resolution technique, *Phys. Rev. X* **3**, 011018 (2013).
- [25] A. Agarwal, H. Wang, C. Schütte, and L. D. Site, Chemical potential of liquids and mixtures via adaptive resolution simulation, *J. Chem. Phys.* **141**, 034102 (2014).
- [26] L. Delle Site, C. Krekeler, J. Whittaker, A. Agarwal, R. Klein, and F. Höfling, Molecular dynamics of open systems: Construction of a mean-field particle reservoir, *Adv. Theory Simul.* **2**, 1900014 (2019).
- [27] A. Gholami, F. Höfling, R. Klein, and L. Delle Site, Thermodynamic relations at the coupling boundary in adaptive resolution simulations for open systems, *Adv. Theory Simul.* **4**, 2000303 (2021).
- [28] K. B. Daly, J. B. Benziger, P. G. Debenedetti, and A. Z. Panagiotopoulos, Massively parallel chemical potential calculation on graphics processing units, *Comput. Phys. Commun.* **183**, 2054 (2012).
- [29] M. Zhao, A. A. Kognole, S. Jo, A. Tao, A. Hazel, and A. D. MacKerell, GPU-specific algorithms for improved solute sampling in grand canonical Monte Carlo simulations, *J. Comput. Chem.* **44**, 1719 (2023).
- [30] K. Binder, B. J. Block, P. Virnau, and A. Tröster, Beyond the van der Waals loop: What can be learned from simulating Lennard-Jones fluids inside the region of phase coexistence, *Am. J. Phys.* **80**, 1099 (2012).
- [31] J. I. Siepmann, I. R. McDonald, and D. Frenkel, Finite-size corrections to the chemical potential, *J. Phys. Condens. Matter* **4**, 679 (1992).
- [32] U. Heinbuch and J. Fischer, On the application of Widom's test particle method to homogeneous and inhomogeneous fluids, *Mol. Simul.* **1**, 109 (1987).
- [33] C. Perego, F. Giberti, and M. Parrinello, Chemical potential calculations in dense liquids using metadynamics, *Eur. Phys. J. Spec. Top.* **225**, 1621 (2016).
- [34] C. Perego, O. Valsson, and M. Parrinello, Chemical potential calculations in non-homogeneous liquids, *J. Chem. Phys.* **149**, 072305 (2018).
- [35] M. Heidari, K. Kremer, R. Cortes-Huerto, and R. Potestio, Spatially resolved thermodynamic integration: An efficient method to compute chemical potentials of dense fluids, *J. Chem. Theory Comput.* **14**, 3409 (2018).
- [36] E. A. Ustinov, Efficient chemical potential evaluation with kinetic Monte Carlo method and non-uniform external potential: Lennard-Jones fluid, liquid, and solid, *J. Chem. Phys.* **147**, 014105 (2017).
- [37] P. S. Clegg, Characterising soft matter using machine learning, *Soft Matter* **17**, 3991 (2021).
- [38] M. Dijkstra and E. Luijten, From predictive modelling to machine learning and reverse engineering of colloidal self-assembly, *Nat. Mater.* **20**, 762 (2021).
- [39] E. Boattini, M. Dijkstra, and L. Filion, Unsupervised learning for local structure detection in colloidal systems, *J. Chem. Phys.* **151**, 154901 (2019).
- [40] G. Campos-Villalobos, E. Boattini, L. Filion, and M. Dijkstra, Machine learning many-body potentials for colloidal systems, *J. Chem. Phys.* **155**, 174902 (2021).
- [41] G. Campos-Villalobos, G. Giunta, S. Marín-Aguilar, and M. Dijkstra, Machine-learning effective many-body potentials for anisotropic particles using orientation-dependent symmetry functions, *J. Chem. Phys.* **157**, 024902 (2022).
- [42] T. Santos-Silva, P. I. C. Teixeira, C. Anquetil-Deck, and D. J. Cleaver, Neural-network approach to modeling liquid crystals in complex confinement, *Phys. Rev. E* **89**, 053316 (2014).
- [43] S.-C. Lin and M. Oettel, A classical density functional from machine learning and a convolutional neural network, *SciPost Phys.* **6**, 025 (2019).
- [44] S.-C. Lin, G. Martius, and M. Oettel, Analytical classical density functionals from an equation learning network, *J. Chem. Phys.* **152**, 021102 (2020).
- [45] P. Cats, S. Kuipers, S. De Wind, R. Van Damme, G. M. Coli, M. Dijkstra, and R. Van Roij, Machine-learning free-energy functionals using density profiles from simulations, *APL Mater.* **9**, 031109 (2021).
- [46] C. Qiao, X. Yu, X. Song, T. Zhao, X. Xu, S. Zhao, and K. E. Gubbins, Enhancing gas solubility in nanopores: A

- combined study using classical density functional theory and machine learning, *Langmuir* **36**, 8527 (2020).
- [47] P. Yatsyshin, S. Kalliadasis, and A. B. Duncan, Physics-constrained Bayesian inference of state functions in classical density-functional theory, *J. Chem. Phys.* **156**, 074105 (2022).
- [48] A. Malpica-Morales, P. Yatsyshin, M. A. Durán-Olivencia, and S. Kalliadasis, Physics-informed Bayesian inference of external potentials in classical density-functional theory, *J. Chem. Phys.* **159**, 104109 (2023).
- [49] X. Fang, M. Gu, and J. Wu, Reliable emulation of complex functionals by active learning with error control, *J. Chem. Phys.* **157**, 214109 (2022).
- [50] A. Simon, J. Weimar, G. Martius, and M. Oettel, Machine learning of a density functional for anisotropic patchy particles, *J. Chem. Theory Comput.* **20**, 1062 (2024).
- [51] A. Simon, L. Belloni, D. Borgis, and M. Oettel, The orientational structure of a model patchy particle fluid: Simulations, integral equations, density functional theory, and machine learning, *J. Chem. Phys.* **162**, 034503 (2025).
- [52] D. de las Heras, T. Zimmermann, F. Sammüller, S. Hermann, and M. Schmidt, Perspective: How to overcome dynamical density functional theory, *J. Phys. Condens. Matter* **35**, 271501 (2023).
- [53] T. Zimmermann, F. Sammüller, S. Hermann, M. Schmidt, and D. de las Heras, Neural force functional for non-equilibrium many-body colloidal systems, *Mach. Learn.: Sci. Technol.* **5**, 035062 (2024).
- [54] F. Sammüller, S. Hermann, D. de las Heras, and M. Schmidt, Neural functional theory for inhomogeneous fluids: Fundamentals and applications, *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2312484120 (2023).
- [55] F. Sammüller, Neural functional theory for inhomogeneous fluids—Tutorial (2023), <https://github.com/sfalmo/NeuralDFT-Tutorial>.
- [56] F. Sammüller, S. Hermann, and M. Schmidt, Why neural functionals suit statistical mechanics, *J. Phys. Condens. Matter* **36**, 243002 (2024).
- [57] F. Sammüller and M. Schmidt, Neural density functionals: Local learning and pair-correlation matching, *Phys. Rev. E* **110**, L032601 (2024).
- [58] F. Sammüller, S. Robitschko, S. Hermann, and M. Schmidt, Hyperdensity functional theory of soft matter, *Phys. Rev. Lett.* **133**, 098201 (2024).
- [59] F. Sammüller and M. Schmidt, Why hyperdensity functionals describe any equilibrium observable, *J. Phys. Condens. Matter* **37**, 083001 (2025).
- [60] S. M. Kampa, F. Sammüller, M. Schmidt, and R. Evans, Metadensity functional theory for classical fluids: Extracting the pair potential, *Phys. Rev. Lett.* **134**, 107301 (2025).
- [61] F. Sammüller, M. Schmidt, and R. Evans, Neural density functional theory of liquid-gas phase coexistence, *Phys. Rev. X* **15**, 011013 (2025).
- [62] M. Buchanan, Machine learning predicts liquid-gas transition, *Physics* **18**, 17 (2025).
- [63] J. Yang, R. Pan, J. Sun, and J. Wu, High-dimensional operator learning for molecular density functional theory, *J. Chem. Theory Comput.* **21**, 5905 (2025).
- [64] A. T. Bui and S. J. Cox, Learning classical density functionals for ionic fluids, *Phys. Rev. Lett.* **134**, 148001 (2025).
- [65] A. T. Bui and S. J. Cox, A first-principles approach to electromechanics in liquids, *J. Phys. Condens. Matter* **37**, 285101 (2025).
- [66] A. T. Bui and S. J. Cox, Dielectrocapillarity for exquisite control of fluids, [arXiv:2503.09855](https://arxiv.org/abs/2503.09855).
- [67] F. Glitsch, J. Weimar, and M. Oettel, Neural density functional theory in higher dimensions with convolutional layers, *Phys. Rev. E* **111**, 055305 (2025).
- [68] J. Dijkman, M. Dijkstra, R. Van Roij, M. Welling, J.-W. Van De Meent, and B. Ensing, Learning neural free-energy functionals with pair-correlation matching, *Phys. Rev. Lett.* **134**, 056103 (2025).
- [69] M. M. Kelley, J. Quinton, K. Fazel, N. Karimitari, C. Sutton, and R. Sundararaman, Bridging electronic and classical density-functional theory using universal machine-learned functional approximations, *J. Chem. Phys.* **161**, 144101 (2024).
- [70] R. Pan, X. Fang, K. Azizzadenesheli, M. Liu-Schiaffini, M. Gu, and J. Wu, Neural operators for forward and inverse potential-density mappings in classical density functional theory, *J. Chem. Phys.* **163**, 164120 (2025).
- [71] 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 (1979).
- [72] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
- [73] R. Pederson, B. Kalita, and K. Burke, Machine learning and density functional theory, *Nat. Rev. Phys.* **4**, 357 (2022).
- [74] R. Stierle, G. Bauer, N. Thiele, B. Bursik, P. Rehner, and J. Gross, Classical density functional theory in three dimensions with GPU-accelerated automatic differentiation: Computational performance analysis using the example of adsorption in covalent-organic frameworks, *Chem. Eng. Sci.* **298**, 120380 (2024).
- [75] The existence of universal density functionals has originally been demonstrated in seminal work for ground states of quantum many-body systems [72]. This formal structure holds analogously for the case of classical fluids [71].
- [76] D. Borgis, R. Assaraf, B. Rotenberg, and R. Vuilleumier, Computation of pair distribution functions and three-dimensional densities with a reduced variance principle, *Mol. Phys.* **111**, 3486 (2013).
- [77] D. de las Heras and M. Schmidt, Better than counting: Density profiles from force sampling, *Phys. Rev. Lett.* **120**, 218001 (2018).
- [78] B. Rotenberg, Use the force! Reduced variance estimators for densities, radial distribution functions, and local mobilities in molecular simulations, *J. Chem. Phys.* **153**, 150902 (2020).
- [79] Code, data, and models available at <https://github.com/sfalmo/neural-mu>.
- [80] J. K. Percus, Equilibrium state of a classical fluid of hard rods in an external field, *J. Stat. Phys.* **15**, 505 (1976).
- [81] D. L. Mobley and J. P. Guthrie, FreeSolv: A database of experimental and calculated hydration free energies, with input files, *J. Comput. Aided Mol. Des.* **28**, 711 (2014).

- [82] S. Luukkonen, L. Belloni, D. Borgis, and M. Levesque, Predicting hydration free energies of the FreeSolv database of drug-like molecules with molecular density functional theory, *J. Chem. Inf. Model.* **60**, 3558 (2020).
- [83] V.P. Sergiievskiy, G. Jeanmairat, M. Levesque, and D. Borgis, Fast computation of solvation free energies with molecular density functional theory: Thermodynamic-ensemble partial molar volume corrections, *J. Phys. Chem. Lett.* **5**, 1935 (2014).
- [84] E.L. Ratkova, D.S. Palmer, and M.V. Fedorov, Solvation thermodynamics of organic molecules by the molecular integral equation theory: Approaching chemical accuracy, *Chem. Rev.* **115**, 6312 (2015).
- [85] K. Binder, J. Horbach, W. Kob, W. Paul, and F. Varnik, Molecular dynamics simulations, *J. Phys. Condens. Matter* **16**, S429 (2004).
- [86] D. de las Heras and M. Schmidt, Full canonical information from grand-potential density-functional theory, *Phys. Rev. Lett.* **113**, 238304 (2014).
- [87] H. Hansen-Goos and R. Roth, Density functional theory for hard-sphere mixtures: The White Bear version mark II, *J. Phys. Condens. Matter* **18**, 8413 (2006).

End Matter

Appendix A: Density functional learning—In the following, we give an overview of machine-learned density functionals, focusing, in particular, on a local learning scheme for the one-body direct correlation functional [54–57,61]. The key element of density functional learning is a neural network that represents a specific density functional relationship. The one-body direct correlation functional $c_1(\mathbf{r}; [\rho], T) = -\delta\beta F_{\text{exc}}([\rho], T)/\delta\rho(\mathbf{r})$ thereby arises as the arguably most natural choice, and it is defined formally as the (scaled) functional derivative of the excess free energy functional $F_{\text{exc}}([\rho], T)$. The neural functional theory by Sammüller *et al.* [54–56,61] rests on representing $c_1(\mathbf{r}; [\rho], T)$ locally by a standard multilayer perceptron, i.e., a fully connected feedforward neural network with multiple hidden layers. Alternatively, it has been proven feasible to represent $F_{\text{exc}}([\rho], T)$ via various neural network architectures [57,68] and to obtain $c_1(\mathbf{r}; [\rho], T)$ in applications from automatic differentiation. Neural density functionals have been shown to facilitate robust and systematic predictions of thermodynamic and structural properties for both homogeneous and heterogeneous fluids in multiscale settings and for states that have not been featured explicitly in the training data [54,56,61].

Choosing $c_1(\mathbf{r}; [\rho], T)$ as the target quantity is motivated by the physical insight that direct correlation functions remain short-ranged for short-ranged interparticle interactions [4], such that a neural network may be constructed to represent the density functional relationship locally: For given position \mathbf{r} , only density information from positions \mathbf{r}' within a narrow window $|\mathbf{r} - \mathbf{r}'| < r_w$ is used to generate the value $c_1(\mathbf{r}; [\rho], T)$. Hence, the neural network mirrors the intended input-output mapping

$$\rho(\mathbf{r}')|_{|\mathbf{r}-\mathbf{r}'|<r_w}, T \rightarrow c_1(\mathbf{r}). \quad (\text{A1})$$

It is typically sufficient to set the spatial cutoff r_w to values of the order of the particle size σ , such as, e.g., $r_w = 2.56\sigma$ [54] for hard spheres of diameter σ and $r_w = 3.5\sigma$ [57,61] for the truncated Lennard-Jones fluid. Local learning has multiple benefits [54–57,61], including the decoupling of predictions from the original simulation box size, which

facilitates the subsequent utilization of the trained neural functional in multiscale applications. We note that reformulations of the local learning scheme in terms of convolutional neural networks are possible and that such approaches have recently been applied to the two-dimensional hard disk fluid [67].

Appendix B: Hard spheres and isothermal Lennard-Jones fluid—We consider in the following the hard sphere fluid with particle diameter σ and the truncated ($r_c = 2.5\sigma$) Lennard-Jones fluid at constant temperature $k_B T = 1.5\epsilon$ and attempt to recover chemical potential values across corresponding datasets via the presented method. As reference, we use grand canonical Monte Carlo simulations with planar inhomogeneous external potentials $V_{\text{ext}}^{(k)}(x)$ and known chemical potentials μ_k . Analogously to the main text, the latter are, however, not prescribed and instead determined during training. The total number of simulations is $n = 450$ for the athermal hard sphere fluid [54] and $n = 500$ for the isothermal Lennard-Jones system [57]. For both fluid types, minimization of the loss function (3) recovers accurately the chemical potential values of the entire datasets. Notably, each resulting neural functional $c_1(x; [\rho])$ delivers the same level of accuracy as the corresponding one that has been trained with prescribed values of μ_k according to Eq. (2), which we recall requires full grand canonical information [54,57]. The chemical potential values and neural network predictions are assessed in Fig. 4. See Ref. [54] for a comparison of neural functional results to analytical treatments [87].

Appendix C: Reduced dataset—We investigate the accuracy of the chemical potential predictions when the dataset size is decreased. Figure 5 shows a comparison for the isothermal Lennard-Jones system, as analyzed for the full dataset of $n = 500$ simulations in Appendix B and Fig. 4(b). Reduced dataset sizes of $n = 100$, 50, and 20 are considered by including only a subset of all simulation results in the training. For $n = 100$ individual simulations, all their chemical potential values can still be recovered reliably. Minor deviations to the reference

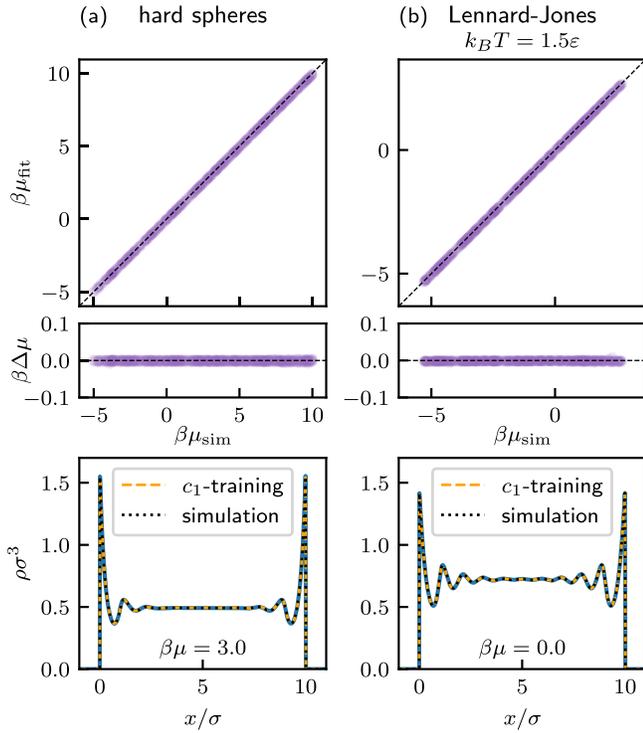


FIG. 4. Results for the hard sphere fluid (a) and the truncated Lennard-Jones fluid at temperature $k_B T = 1.5\epsilon$ (b) analogously to Fig. 2. The values $\beta\mu_{\text{fit}}$ recovered via the present procedure agree with the simulation reference $\beta\mu_{\text{sim}}$ (top row), as confirmed by the small values of the scaled potential difference $\beta\Delta\mu = \beta\mu_{\text{fit}} - \beta\mu_{\text{sim}}$ (middle row). The density profiles (bottom row, solid blue lines) correspond to confinement between hard walls of separation distance 10σ and are obtained via self-consistent solution of Eq. (1) with the neural functional $c_1(x; [\rho])$ resulting from the minimization of the loss function (3). For comparison, we show results of neural functionals that have been trained directly with grand canonical data using known values of μ (dashed orange lines) as well as density profiles from grand canonical Monte Carlo simulations (dotted black lines).

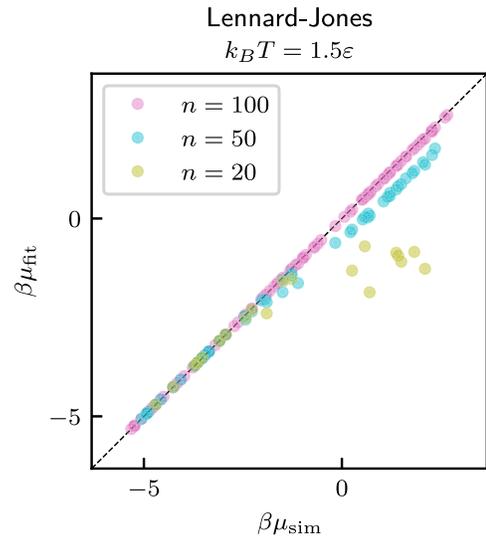


FIG. 5. Comparison of results for training with reduced dataset sizes $n = 100, 50,$ and 20 , as indicated. Scaled chemical potential values $\beta\mu_{\text{fit}}$ are recovered reliably for the case $n = 100$. Deviations to the reference data $\beta\mu_{\text{sim}}$ become noticeable for $n = 50$. For $n = 20$, the procedure yields correct results only for small chemical potentials, while larger values are underestimated.

values start to appear, particularly for larger chemical potentials, for $n = 50$. If only $n = 20$ simulations are considered, recovering chemical potentials remains possible for small values of μ , but the quality of the results worsens for larger chemical potential values, which tend to be underestimated. This behavior is rationalized by recalling the underlying mechanism of the procedure, which hinges on the successful determination of the functional mapping $c_1(\mathbf{r}; [\rho], T)$. For large values of μ , the fluid is dense and highly correlated, such that it becomes particularly intricate to capture accurately the density functional dependence in this regime with only few training samples.