RESEARCH ARTICLE | SEPTEMBER 08 2025

Rethinking the Kohn–Sham inverse problem

Alexander Kaiser; Stephan Kümmel

J. Chem. Phys. 163, 104101 (2025)
https://doi.org/10.1063/5.0281993

The Journal of Chemical Physics ARTICLE pubs.aip.org/aip/jcp

Rethinking the Kohn–Sham inverse problem

Cite as: J. Chem. Phys. 163, 104101 (2025); doi: 10.1063/5.0281993
Submitted: 22 May 2025 • Accepted: 21 August 2025 •
Published Online: 8 September 2025

Alexander Kaiser and Stephan Kümmel[a)]

AFFILIATIONS

Theoretical Physics IV, University of Bayreuth, 95447 Bayreuth, Germany

[a)]Author to whom correspondence should be addressed: stephan.kuemmel@uni-bayreuth.de

ABSTRACT

Density functional theory (DFT) is a cornerstone of modern electronic structure theory. In the Kohn–Sham scheme, the many-electron Schrödinger equation is replaced by a set of effective single-particle equations. Thus, the full complexity of the quantum mechanical many-particle effects is mapped to the exchange–correlation potential $v_{\mathrm{xc}}(\mathbf{r})$. Almost all DFT calculations done in practice rely on approximations to $v_{\mathrm{xc}}(\mathbf{r})$. However, numerical representations of the quasi-exact $v_{\mathrm{xc}}(\mathbf{r})$ can be obtained from quasi-exact densities by inverting the Kohn–Sham procedure. This inverse Kohn–Sham scheme is an important source of insight into exact DFT. Here, we review the inverse Kohn–Sham problem and explain in detail several aspects of why Kohn–Sham inversion is intrinsically difficult. We then present several inversion schemes and discuss their pros and cons, specifically addressing the effects of statistical uncertainties that are inevitable in quantum Monte Carlo reference densities. We use these schemes to obtain representations of $v_{\mathrm{xc}}(\mathbf{r})$ that correspond to the ground-state densities that have become available from accurate diffusion Monte Carlo calculations on real space grids for the Li_2 and N_2 molecules, and the C atom. In the latter, the highest occupied orbital has a nodal line and the exchange–correlation potential goes to a different asymptotic value in this direction. As an outlook, we discuss the possibility of interlacing quantum Monte Carlo and Kohn–Sham theory by using the quasi-exact Kohn–Sham determinant to fix the nodes in a diffusion Monte Carlo calculation.

I. INTRODUCTION

DFT is one of the most frequently used theories for calculating the electronic structure of molecules and solids. The accuracy of any DFT calculation depends decisively on the approximations that are chosen for the exchange (x) and correlation (c) functionals E_{xc} and v_{xc}, respectively. While the existence of the exact E_{xc} and v_{xc} is guaranteed, and while many successful approximations have been developed, gaining insight into the properties of the *exact* functionals continues to be a challenge as their general form is not explicitly known. Inversion, i.e., constructing the numerically quasi-exact v_{xc} from a very accurate density that has been obtained by some other means, typically an accurate wavefunction calculation, is an important way of obtaining insight into the properties of the exact v_{xc}.[1–27] As the true inverse I is nothing else than the functional derivative of the true xc energy functional $I = \delta E_{\mathrm{xc}}/\delta n$, inversion is at the same time an indirect way of revealing properties of the true E_{xc}. Therefore, inverting the Kohn–Sham problem is of considerable interest in DFT.

Furthermore (allowing ourselves a brave look into the future), if one would have inversion schemes that were so reliable that they could be automated, then such schemes could pave the way for machine learning of the Hohenberg–Kohn map or more precisely the map between the true density and the auxiliary Kohn–Sham (xc) potential: with a reasonable number of quasi-exact densities that exist from calculations with various wave function based methods, and their inverted xc potentials, sufficient training data might be available to machine-learn the potential-density map. From the inverted density, one would have access to the exact orbitals and other quantities such as the non-interacting kinetic energy density τ. With those variables, one could locally estimate the xc potential, e.g., from a machine-learned function $v_{\mathrm{xc}}(n_\sigma, \nabla n_\sigma, \tau)$ in some of the variable parameterizations that are common for meta-GGAs. This could complement the learning of ϵ_{xc}, cf. Refs. 28 and 29.

In this article, we revisit the Kohn–Sham inversion problem. We first discuss it on general grounds, then discuss the pros and cons of algorithms that can be used for inversions, and finally, present quasi-exact v_{xc} potentials of molecules that were obtained from inverting the densities from accurate Quantum Monte Carlo calculations. In an outlook, we discuss the possibility of combining Kohn–Sham inversion with diffusion Monte Carlo (DMC) calculations.

II. THE KOHN–SHAM MAP

The Hohenberg–Kohn theorem[30] (HKT) establishes a one-to-one correspondence between density n and potential v up to a constant. Therefore, any Kohn–Sham potential directly leads to a unique density and, in principle, any density must also lead to a unique (up to a constant) Kohn–Sham potential. For the sake of clarity, we omit spin indices here and in the following whenever it is reasonable, even though some expressions might genuinely be spin-resolved. Since the Hartree and the external potential (given the particle number) have their own one-to-one mapping with the density in finite systems, the potential-density-map not only connects the Kohn–Sham potential with the density but also with the xc potential itself. Thus, we use the symbols for forward and inverse maps for the Kohn–Sham as well as the xc potentials. The (forward) map from potential to density is numerically rather simple, requiring the solution of the Kohn–Sham equations for all orbitals contributing to the density,

$$F[v_{\mathrm{xc}}] = n. \tag{1}$$

The inverse map,

$$I[n] = v_{\mathrm{xc}}, \tag{2}$$

although guaranteed to exist and be unique by the HKT, is a much more intricate and difficult to realize mapping. It belongs into the class of inverse problems[31] and is ill-posed,[32] where well-posedness requires the existence of a solution, its uniqueness, and the continuous dependence on the initial conditions. A general approach to inverse problems is based on regularization strategies,[31] e.g., some well-chosen transformation of the input data (in our case the density), and restrictions on the possible outcomes (the xc potential), to arrive at a well-posed problem or at least render the inverse problem less ill-posed. With too little regularization, the ill-posedness is still present, while too much regularization puts an unnecessary bias onto the inverted potential. Between these limiting cases, there usually is an optimal range for the regularization and often much effort has to be spent to find it. Determining the corresponding potential for a given density (or given spin-densities) is a challenging task. Conversely, the quality of any candidate potential can readily be evaluated by its density's difference to the target density since the forward map is realized with high accuracy at low computational cost. The HKT guarantees the one-to-one correspondence mathematically, based on the variational principle by an energy difference. The classical proof[30] of the HKT exploits that analytically one can identify an infinitesimal energy difference. In the numerical picture, the guaranteed one-to-one correspondence becomes volatile to the point where it breaks down because all real numerics work with a limited precision. A simple, well-known example for the non-uniqueness of the potential in a finite-precision numerical realization is the following: take any molecular potential-density pair and add any positive function of small spatial extent to the original potential in the far asymptotic region. The additional repulsive potential does not alter the density in a numerically noticeable way because in the far asymptotic region, the density already has decayed to a number that is negligible in a calculation with finite numerical precision. Thus, the density corresponding to the manipulated potential will be identical to the original density within numerical accuracy, disproving the uniqueness of the potential to density map in a numerical representation. Analytically, however, the manipulated potential is linked to a new density that differs from the original density, even if only ever so slightly.

Revisiting the proof of the HKT with the intention to allow for differing (spin-)densities and Kohn–Sham Hamiltonians yields insight into the degree of uniqueness that the HKT guarantees. With this strong form of the HKT,[33–35] our aim is to probe variations of a given potential and the corresponding variations of the density to make the consequences of differing potentials with similar densities more transparent. For any two Kohn–Sham potentials that share the same external potential, $v_{\mathrm{KS}\sigma}$ and $v'_{\mathrm{KS}\sigma}$, and their corresponding densities, n_σ and n'_σ, the argument of the HKT is based on the non-positive quantity,

$$J = \sum_\sigma \int \left(v_{\mathrm{KS}\sigma} - v'_{\mathrm{KS}\sigma}\right)\left(n_\sigma - n'_\sigma\right) \mathrm{d}^3 r \leq 0. \tag{3}$$

Here and in the following, integrals without explicitly noted boundaries denote volume integrations over the functions' full domain. If the two potentials are not equal, $v_{\mathrm{KS}\sigma} \neq v'_{\mathrm{KS}\sigma}$, then J has to be strictly negative, $J < 0$, due to the variational principle.[30,33] Identical densities would contradict a negative J, hence $n_\sigma \neq n'_\sigma$, uniqueness is guaranteed, and the HKT follows. When the density (or potential) of one density-potential pair approaches the density (or potential) of the other pair, the absolute value of J decreases until it vanishes in the limiting case of identical densities (or potentials up to a constant). J is at least quadratic in the density-potential variation since a linear variation in density (or potential) induces a linear variation in the potential (or density). Equation (3) allows quantifying the limitations on the numerical realization of the uniqueness of the potential-density map. While the proof of the HKT starts from a given potential, J treats density and potential equally and allows assessing the consequences of small differences in the density. With the numerical resolution restricted to a small positive number ϵ, representing, e.g., the floating point precision, one has $J < -\epsilon$. In this line of thought, any two numerical densities with $J \in [-\epsilon, 0]$ cannot be distinguished from each other. While ϵ might be very small, the exponentially decaying asymptotic tail of the density difference in Eq. (3) allows for significant variations in the potentials.

Splitting the Kohn–Sham potential in is constituents, allowing the external potential to cancel, and carrying out spin summation in the Hartree part, the fraction of J corresponding to the xc potential is extracted,

$$\begin{aligned} J &= \int \left(v_{\mathrm{H}} - v'_{\mathrm{H}}\right)\left(n - n'\right) \mathrm{d}^3 r + \sum_\sigma \int \left(v_{\mathrm{xc}\sigma} - v'_{\mathrm{xc}\sigma}\right)\left(n_\sigma - n'_\sigma\right) \mathrm{d}^3 r \\ &= J_{\mathrm{H}} + \sum_\sigma J_{\mathrm{xc}\sigma}. \end{aligned} \tag{4}$$

The Hartree contribution is a positive quantity $J_{\mathrm{H}} > 0$, which can be seen by the relation $-4\pi n = \nabla^2 v_{\mathrm{H}}$ and the first Green identity. J_{H} plays a role in many inversion methods and has repeatedly been identified[3,36–38] as a relevant quantity. The above-mentioned arguments also transfer to J_{xc} and show that xc potentials for densities that agree within numerical resolution can differ significantly, e.g., in regions of low density.

A weaker form of ambiguity of potential-to-density map stems from the Coulomb singularity at the nuclei. In the presence of a

$Z/|\mathbf{r}-\mathbf{R}|$ divergence for a nucleus at $\mathbf{R}$ with charge Z, the value of the xc potential at the exact location of the nucleus becomes irrelevant. Although those points build a null set, the Coulomb divergence dominates in the close vicinity of the nuclei and renders the values of the xc potential less important in those regions. Thus, the forward map is rather insensitive to variations of the potential at the nuclei. For example, the xc potential associated with generalized gradient approximations (GGAs) typically diverges at the nuclei, yet GGAs can successfully predict energies and other observables.

However, the same consideration for the inverse map yields a very different picture: a small variation of the density at the nuclei typically results in a very large variation in the potential. In order for the density to correspond to a molecular Hamiltonian, it has to satisfy the cusp condition,

$$\frac{\mathrm{d}}{\mathrm{d}|\boldsymbol{\epsilon}|}\ln \hat{n}(|\boldsymbol{\epsilon}|) = -2Z, \tag{5}$$

with a separation vector $\boldsymbol{\epsilon} = \mathbf{r} - \mathbf{R}$ and $|\boldsymbol{\epsilon}|$ being the distance between the nucleus and some coordinate $\mathbf{r}$, where $\hat{n}$ is the density spherically averaged with respect to the nucleus' position. This fixes the density to an exponentially decaying form, with an exponent proportional to Z. If the density does obey the cusp condition, but with a slightly different exponent $\tilde{Z} \neq Z$, the xc potential has to account for this deviation, as the external potential has a fixed singularity $Z/|\boldsymbol{\epsilon}|$. As a consequence, the xc potential will develop a singularity of the type $(\tilde{Z}-Z)/|\boldsymbol{\epsilon}|$. Thus, even for tiny deviations $\tilde{Z} - Z \neq 0$, the spatial form of the inverted xc potential will be massively affected: small variations of the density yield diverging differences in the potential. The inversion is thus ill-posed: while the forward map is straightforward, the inverse map is sensitive to small changes in the density and many candidate potentials produce densities very close to the reference density.

Numerically realizing an inversion, therefore, requires one to handle the ambiguity in the potentials as good as possible. With ambiguity, there is no predetermination of some characteristics of the xc potential. This seemingly unfavorable aspect of inversion can be reinterpreted as an advantage because it ultimately allows the user to select those characteristics. This is part of the regularization that is necessary to handle inverse problems. Since the regularization influences the inversion outcome, inversion schemes have to be chosen carefully. Figure 1 illustrates how different inversion methods that start from the same density can yield different inverted potentials. For example, the asymptotic behavior of the xc potential is known to be $-1/|\mathbf{r}|$ from analytic considerations;[39,40] therefore, this exact property can be incorporated into the inverted potential manually. This is often realized with the Fermi–Amaldi potential,[23,24]

$$v_{\mathrm{xc}}^{\mathrm{FA}} = -N^{-1} v_{\mathrm{H}}, \tag{6}$$

or other prototype xc potentials with correct asymptotic decay. This example paradigmatically reveals pros and cons of regularization strategies: For many systems, the Fermi–Amaldi potential imposes the proper far-asymptotic structure in a stable and efficient way. However, for systems in which the highest occupied orbital has a nodal surface, the xc potential goes to different asymptotic constants in different directions[41–43] and this feature is not captured by the Fermi–Amaldi potential. Thus, we see here an example for the fact

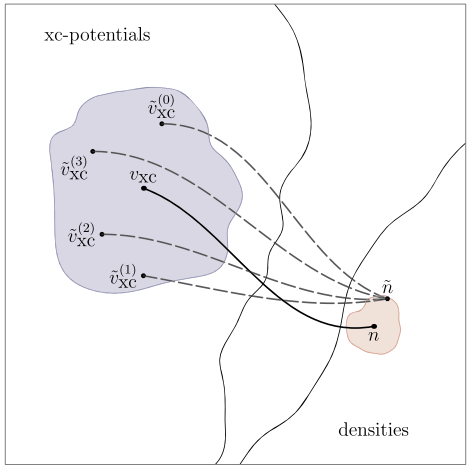

FIG. 1. Ambiguity: a schematic sketch of the mapping from xc potential space to density space. The true map from the potential v_{xc} to density n is indicated by the solid line. Any inversion (dashed lines) starts from the numerical approximation or representation $\tilde{n}$ of the true density n. Using different inversion methods, the approximate density $\tilde{n}$ maps back to different inverted potentials $\tilde{v}_{\mathrm{xc}}^{(0)}, \ldots \tilde{v}_{\mathrm{xc}}^{(3)}$. While the differences in the density domain might appear small, the differences between the inverted potentials can be large. In other words, a "small basin" in density space can map to a "large basin" in potential space. Conversely, distinct potentials can result in visually very similar densities. The question arises how to measure or quantify the distinctness of densities in contrast to the distinctness of the corresponding potentials.

that regularization techniques might need to be chosen suitably and specifically for the system of interest.

The ambiguity in the domain of the potentials, on the one hand, simplifies the development of E_{xc} functionals: while tailored variations in the xc energy functional result in desired variations in the energy (and less considered variations in the xc potential), the density might remain almost unchanged. This offers a path to lower the functional driven error[44] of a density functional approximation, without having to worry too much whether one might increase the density-driven error. On the other hand, it also imposes a challenge: there is no universal, straightforward strategy of E_{xc}-construction for improving cases in which density-driven errors are the limiting factor. Therefore, such strategies have to be developed specifically for each class of functionals, as successfully exemplified in recent years for, e.g., meta-GGAS[45,46] and range-separated hybrids.[47,48]

With a substantial part of the mathematical rigor of the HKT hidden in numerically inaccessible features of the density, numerical realization of the inverse may be beyond reach in many cases. With these limitations in mind, let us examine the use case of inversions, namely, the generation of xc potentials from high-quality densities. Such densities most of the time originate from a method other than DFT, i.e., coupled cluster, Møller–Plesset perturbation theory, configuration interaction, Quantum Monte Carlo methods, or newly emerging machine-learning techniques,[49] to only name a few. Those densities might be very reasonable approximations to the true density, but are rarely ever exactly equal to the true density for two reasons. In any of the methods, there remains some conceptual convergence error. This can be, e.g., the truncation of an expansion

in Slater determinants or the fixed node approximation in DMC. Another error is given by the numerical representation, most often due but not limited to a finite basis set. Any inverted xc potential can only refer to the (highly accurate but approximate) numerical density, which will be called the reference density from here on. While the errors in the reference densities might be negligible for almost all applications, inverse problems are pathologically susceptible to numerical errors of the input data,[31] i.e., the reference density. We will discuss this in detail for the case of the statistical fluctuations in DMC densities. In order to reproduce the slight inaccuracies of the reference density, the inverted potential tends to build up spurious features, in some cases to an extent where the inverted potential is useless.

In addition to the inverse problem being ill-posed, the v-representability problem can also prevent one from being able to solve the inversion problem. Even when the exact interacting density is non-interacting v-representable, its numerical representation might not be. In such a case, all the criteria for well-posedness, i.e., existence, uniqueness, and stability of the solution, might be violated. The same applies if the exact density itself is not non-interacting v-representable. In the following, however, we always assume that the densities that we feed into the inversion are non-interacting v-representable.

III. INVERSION TECHNIQUES

The problem of inverting the Kohn–Sham map has been addressed many times and with various methods. Reference 1 was among the first to propose an iterative scheme that updates an initial potential with feedback from the difference of the current and the reference density. Reference 3 devised an extrapolation scheme to yield an inverted potential in the limit that the Hartree potential of the reference and the inverted density be identical, i.e., $J_H = 0$. The work in Refs. 50 and 51 retrieves the inverted potential from the comparison of a suitable DFT energy density with the corresponding wave function energy density. Many more approaches to the inverse problem exist.[1–11,15,16,19,21–25,52]

In the following, we present iterative schemes to improve some initially guessed xc potential by repeated application of an update rule until convergence. While our reference density is denoted by n, the Kohn–Sham density of the current iteration is denoted by $n^{(i)}$, or n' if the precise number of the iteration step does not matter. We consider update rules that can be written as

$$v_{\mathrm{xc}\sigma}^{(i+1)} = v_{\mathrm{xc}\sigma}^{(i)} + \epsilon \Delta v_{\mathrm{xc}\sigma}[n^{(i)}, n], \tag{7}$$

where ϵ is some feedback strength to manipulate the speed of convergence and $\Delta v_{\mathrm{xc}}[n', n]$ aims to improve the potential with information from the current and reference densities, n' and n, respectively. Thus, one typically works with an iterative scheme of the type,

$$v_{\mathrm{xc}}^{(0)} \mapsto n_\sigma^{(0)} \mapsto \Delta v_{\mathrm{xc}\sigma}[n^{(0)}, n] \mapsto v_{\mathrm{xc}\sigma}^{(1)} \mapsto \ldots. \tag{8}$$

Any iterative update scheme has an intrinsic starting point dependence via the initial xc potential $v_{\mathrm{xc}}^{(0)}$. Two different starting potentials can converge to different self-consistent inverted potentials with an almost identical density, as schematically illustrated in Fig. 2.

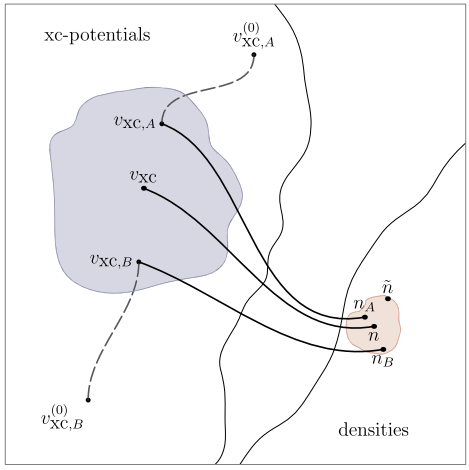

FIG. 2. Starting point dependence: when the inversion is based on iterative potential updates, the initial xc potential can introduce a significant bias into the converged inverted potential. An important example is features in the asymptotic region of the target xc potential corresponding to the reference density. Typically, these cannot be uniquely reconstructed.

Since the inverted potential is almost invariant under changes of the potential in low-density regions, as discussed in Sec. II, the asymptotic behavior of the initial density is practically left unaltered throughout the iterations.

However, solving the inverse problem with an update rule method has an important technical advantage. Most update rules fit into existing DFT codes with little effort, as the update of the xc potential in the inverse problem is computationally very similar to the update of the xc potential in a self-consistent forward problem. This is applicable to update rules that fall into the class of LDAs, GGAs, and even meta-GGAs, where we denote an update rule as "falling into the LDA, GGA, or meta-GGA class" when the update procedure requires only knowledge of the density, or density and gradient, or in addition also the kinetic energy density, respectively.

The decisive quantity of the update rule is the update functional $\Delta v_{\mathrm{xc}}[n', n]$. It should provide corrections to the potential so that the current density n' is changed in such a way that it gets closer to the reference density n. The optimal update functional is the inverse response function χ^{-1},

$$\delta v(\mathbf{r}) = \int \chi^{-1}(\mathbf{r}, \mathbf{r}') \delta n(\mathbf{r}') \, \mathrm{d}^3 r', \tag{9}$$

which, however, is generally not known as an explicit functional of the density. A successful update functional needs to satisfy several challenging requirements. First, the density spans a wide range of values, often over several orders of magnitude, while the xc potential has a very limited range of values (confined in the interval $[-10, 1]$ a.u. for the potentials found in this work). Second, the update functional has to convert the input densities to physically reasonable potential updates. For efficient updates, this requires the update functional to have some incorporated knowledge of the non-locality of the Kohn–Sham equations, i.e., how a variation of the potential alters the density. The more of this knowledge is implemented into an update rule, the better it can predict reasonable potentials.

A common idea[34,53] is the *a priori* determination of the direction of the update $\Delta v_{\mathrm{xc}\sigma}[n', n]$. Imagine that at some point $\mathbf{r}$, the value of the current density $n'(\mathbf{r})$ is less than the value of the reference density $n(\mathbf{r})$. For the densities to align, the update to the current potential should be negative since a more negative potential attracts more density, and thus increases $n'(\mathbf{r})$. If Δv_{xc} has this property globally, i.e., $\Delta v_{\mathrm{xc}}[n', n](\mathbf{r}) < 0$ when $n'(\mathbf{r}) < n(\mathbf{r})$ and vice versa for all $\mathbf{r}$, we call Δv_{xc} sign-consistent. Sign-consistency is not a necessary requirement for an update rule to be successful since the Kohn-Sham equations themselves and the normalization of the density are non-local operations. Thus, even if the update rule is sign-inconsistent, the Kohn–Sham scheme can (by the semi-local Laplacian, the convoluting non-local Hartree potential, and normalization) decrease the density in the next iteration at some point despite the potential being decreased at the same point. Usually (in our experience), the sign-consistency helps accelerate convergence and update rules with a large violation of the sign-consistency tend to give less accurate results.

Many choices for update rules can be found in the literature.[3,19,37,38,50,51,53,54] In the course of the research that led to this paper, we have tested many update rules, but in the following, we focus on a few ones because they appeared as particularly advantageous to us.

To start with, the Hartree potential offers a special access to the inversion procedure. Up to a constant, the Hartree potential has a unique relation to the density since the Poisson equation grants uniqueness for given boundary conditions. If the Hartree potentials of two densities align, the densities must be equal. Furthermore, since the Hartree potential readily converts a density to a physical potential, the update,

$$\Delta v_{\mathrm{xc}\sigma}^{\mathrm{HAR}} = v_{\mathrm{H}}[n'] - v_{\mathrm{H}}[n] \tag{10}$$

is well-known in inversion.

The update is connected to the minimization of J_H from Eq. (4), where J plays the central role of a minimizing functional in the inversion method devised by Zhao, Morrison, and Parr.[3] The Hartree update to the potential is arbitrarily smooth[55] due to the integration of the density, i.e., it does not inherit any possible cusps from numerical artifacts of the reference density. Its non-locality usually accelerates convergence, but simultaneously violates sign-consistency. While the Hartree update is locally sign-consistent in a large portion of the density's domain, at points of rapid density changes, including the nuclear cusps, the smoothing character of the Hartree potential prevents it to resolve sharp details. The Hartree potential at a nuclear cusp is smooth and of quadratic order with the distance to the nucleus.

From a multipole expansion of the Hartree potential, one can infer that the update has the form of a dipole potential in the asymptotic domain, $\Delta v_{\mathrm{xc}\sigma}^{\mathrm{HAR}} \to \mathbf{p}\cdot\hat{\mathbf{r}}/|\mathbf{r}|^2$ with some vector $\mathbf{p}$. Therefore, slight changes in the asymptotic regions are possible with the Hartree update without "overwriting" an already existing $-1/|\mathbf{r}|$ behavior from the initial xc potential. The Hartree potential has no spin dependence as it models electrostatic repulsion from electrons independent of the spin. However, for an update rule in the inversion of a spin-polarized system, utilization of spin-dependent Hartree potentials, $v_{\mathrm{H}_\sigma} = v_{\mathrm{H}}[n_\sigma]$, is advantageous since the density feedback is resolved separately in each spin channel. Furthermore, the one-to-one correspondence due to the Poisson equation only holds per spin channel in the spin polarized case.

Another conceptually very transparent update rule[53] is

$$\Delta v_{\mathrm{xc}\sigma}^{\mathrm{DEN}} = n'_\sigma - n_\sigma. \tag{11}$$

It can be understood as deriving from Eq. (9) under the approximation that the inverse response function is purely local. In Eq. (11), the update to the potential at some point in space stems only from the density difference at this point. This update is sign-consistent and needs to gain any non-local character from repeated updates that include non-locality from the Kohn–Sham equations into the inverted potential. A serious limitation of this update rule is how it handles the asymptotic region: since $\Delta v_{\mathrm{xc}\sigma}^{\mathrm{DEN}}$ decays exponentially, as does the density, the update is unable to alter the potential in the asymptotic and even near-asymptotic regions. Yet, not only in the asymptotic region, but quite generally the update does not map the wide range of density values to moderate xc potential values. This also seriously affects the nuclear region for the following reason: in an all-electron calculation, the density is largest at the nuclei and it usually changes its value at the nuclei significantly from iteration to iteration in a self-consistent calculation, e.g., due to normalization. The density differences in Eq. (11) at the nuclei inherit the large magnitude of these changes and map it to the potential update. This leads to large spikes in the potential at the nuclei, easily overshooting the optimal update. In the subsequent iteration, the overshooting is corrected for with maybe an even larger overshooting in the opposite direction. This results in oscillations in the iterative scheme and often also introduces characteristic wiggles in the potential near the nuclei that may not be restricted to just the close vicinity of the nuclei. To prevent such behavior, the feedback strength ϵ has to be lowered, even if the update would need strong feedback elsewhere, e.g., in the asymptotic domain. Modifications to scale the update by multiplication with powers of the distance, $|\mathbf{r}|(n'_\sigma - n_\sigma)$, in spherical systems, or in general some local weight $w(\mathbf{r})(n'_\sigma - n_\sigma)$, to counteract the imbalance from nuclear to asymptotic density values, are presented in Ref. 53. While these can help improve the update, and while the update can work well for soft-featured model potentials,[10] we have not found sufficient relieve from the above-mentioned disadvantages in the all-electron atomic and molecular calculations that are the focus of our present work.

To address the above-mentioned issues, we introduce a weighting of the density difference with the inverse Hartree potential, $w(\mathbf{r}) = v_{\mathrm{H}}^{-\beta}[n]$. Due to its smoothness and its strong tendency to follow the magnitude of the density, the Hartree potential is an optimal candidate for a weight function, without the necessity for the user to optimize parameters. In regions of high density (and thus a large Hartree potential), absolute density differences $n'_\sigma - n_\sigma$ become large, even if the relative density differences are small. The Hartree-weighting reduces such density differences to smaller values, improving the update at, e.g., the nuclei. For large relative density differences at small absolute difference, e.g., in the asymptotic domain, the weighting instead increases the feedback. One could argue that a weighting that accomplishes the same and is yet simpler is the inverse density $w(\mathbf{r}) = n^{-1}(\mathbf{r})$. However, this weighting becomes problematic in the asymptotic region of finite systems: There, the density decays exponentially, and thus, the potential update $(n'_\sigma - n_\sigma)/n_\sigma$ is likely to eventually diverge exponentially.

The potential update resulting from the "density over Hartree" approach is

$$\Delta v_{\mathrm{xc}\sigma}^{\mathrm{DoH}} = \frac{n'_\sigma - n_\sigma}{v_{\mathrm{H}}^{\beta}[n]}. \tag{12}$$

The scaling parameter β allows for adaptation of the method, affecting the damping at the nuclei and the asymptotic behavior, in analogy to the parameter in Ref. 53. For the purpose of this work, we only present results with $\beta = 1$ for the DoH rule.

Another approach that yields reasonable potential updates is utilizing some density functional approximation with potential $v_{\mathrm{xc}}[n]$. The update then reads $\Delta v_{\mathrm{xc}\sigma} = -\left(v_{\mathrm{xc}\sigma}[n'_\sigma] - v_{\mathrm{xc}\sigma}[n_\sigma]\right)$. For the exact xc potential function, this update would remove the potential of the current density n'_σ and substitute it with the exact target potential $v_{\mathrm{xc}\sigma}[n_\sigma]$. Since the xc potentials of many reasonable density functional approximations try to resemble the exact potential, one can utilize them in the inversion process. This approach also maps density values to corresponding potential values of reasonable magnitude. Unfortunately, many density functionals yield potentials that have spurious features (e.g., GGAs close to the nucleus). Furthermore, it has been argued that some recent developments in DFT have led to better ground-state energies, but did not necessarily lead to improved densities.[56] Therefore, and because it is numerically straightforward, we only use the local density approximation here. It yields the update rule,

$$\Delta v_{\mathrm{xc}\sigma}^{\mathrm{LDA}} = -\left(v_{\mathrm{x}\sigma}^{\mathrm{LDA}}[n'_\sigma] - v_{\mathrm{x}\sigma}^{\mathrm{LDA}}[n_\sigma]\right). \tag{13}$$

For the sake of transparency, we employ only the exchange part of the LDA $v_{\mathrm{x}}^{\mathrm{LDA}} = -A_x n^{\frac{1}{3}}$ in our update rule. The one-third exponent decreases large densities and enhances lower densities and thus is useful for mapping the density scale to the potential scale. While strictly local and sign-consistent, oscillations at the nuclei are still possible, yet typically much smaller than in the DEN update.

Our final update rule combines the inverse Hartree weight with the updates from the LDA with some adjustments. Apparently, the density to the power of one third as in the exchange-only LDA is a good mapping of density values to potential values since in self-consistent LDA calculations, the xc potential shows a range of values in good agreement with the values of the true xc potential. When dividing the LDA potential (differences) by the Hartree potential, this property is lost. Therefore, we introduce exponents at the potentials via a parameter β and tailor the expression such that it yields the LDA update for $\beta = 0$. The resulting, unmodified update essentially combines LDA and Hartree in the form

$$\frac{\left|v_{\mathrm{x}\sigma}^{\mathrm{LDA}}[n'_\sigma]\right|^{(1+\beta)} - \left|v_{\mathrm{x}\sigma}^{\mathrm{LDA}}[n_\sigma]\right|^{(1+\beta)}}{v_{\mathrm{H}}^{\beta}[n]}. \tag{14}$$

The prefactors involved in $v_{\mathrm{x}\sigma}^{\mathrm{LDA}}$, i.e., A_x and the factors related to spin scaling for polarized densities, only introduce a global prefactor for the potential update, which unnecessarily scales with β and forces the user to adapt the update strengths ϵ when varying β. We have found the update rule to be more consistent across different systems, i.e., requiring less system-specific adjustment, without the scaled prefactors. Since the Hartree potential is usually larger than the xc potential, the above-mentioned update has the tendency to show globally lower amplitudes for larger β. We, therefore, also divide the Hartree potential in the denominator by the particle number N. The resulting "LDA over Hartree" update is given by

$$\Delta v_{\mathrm{xc}\sigma}^{\mathrm{LoH}} = \frac{{n'_\sigma}^{(1+\beta)/3} - {n_\sigma}^{(1+\beta)/3}}{N^{-\beta} v_{\mathrm{H}}^{\beta}[n]}. \tag{15}$$

We denote the method by LoH(β) to state the value used for β. Thus, LoH(0) is equivalent to the LDA-based update rule. The LoH(1) update does not directly resemble the DoH update, but it corrects the exponents of the densities in the numerator of Eq. (12) to 2/3 when using a denominator with the non-exponentiated Hartree potential. The LoH rule is sign-consistent. It damps the density feedback at the nuclei and improves the decay of the update in the asymptotic region by weighting and adjusting the density values with a power law. The parameter β can be chosen in a relatively large range without affecting the results too much. At the same time, the dependence on β originates from the interplay of numerator and denominator, and therefore, the effects that changing beta has on the update rule's asymptotic behavior are somewhat subtle and not always straightforward to foresee.

We illustrate the update rules in Fig. 3, where the updates Δv_{xc} in the course of an inversion of the reference xKLI density (with initial LDA potential and the update rule HAR) at iteration 200 are plotted. The updates Δv_{xc} of course differ in absolute magnitude; therefore, we normalize them by division by their largest absolute value. The HAR update rule is not sign-consistent but follows the sign of the density difference roughly. The sign-consistent rules share nodes with the difference of the current and the reference density. The update rules HAR, LDA, and LoH are able to update the inverted potential in the asymptotic regions, contrary to the exponentially decaying DEN and DoH rules. The rule LoH(0), which is equivalent to the LDA update, has a reduced exponential decay (by a factor of 1/3) and gets magnified in the asymptotic region due to the inverse Hartree weight, which renders it capable of adjusting the asymptotic regions. The LoH rules of varying parameters show that small values for β enforce a preferable scaling in the asymptotic region. Larger values of β have fast exponential decay (with values of $\beta > 3$ even faster than for DEN or DoH) and focus on corrections to the regions around the nuclei. On the other hand, low values of β map the orders of magnitude of the density to an increasingly low range of potential values. A reasonable range of β is, therefore, [0,3], while values in the interval $]-1, 0]$ still yield a sign consistent update, but tend to become less efficient as they do not take into account the density any more as $\beta \to -1$. In Fig. S4 of the supplementary material, we show the same graph for a reference density including numerical artifacts.

It is possible and in general also a good idea to combine the above-mentioned update rules instead of using a single one. We use different update rules combined together in one calculation, but sequential application is also possible. Since each rule can use its own feedback strength ϵ, cf. Eq. (7), we are free to assign feedback parameters to the ingredient rules separately. This defines a large space of parameters to explore and we explain which choices we took in the following. Since the Hartree update reaches out to the asymptotic regions, has non-locality included, and has superior smoothness, it is our base ingredient. Experience has shown that a feedback strength of $\epsilon_{\mathrm{HAR}} = 0.5$ is often a good choice. On top of the HAR rule, one

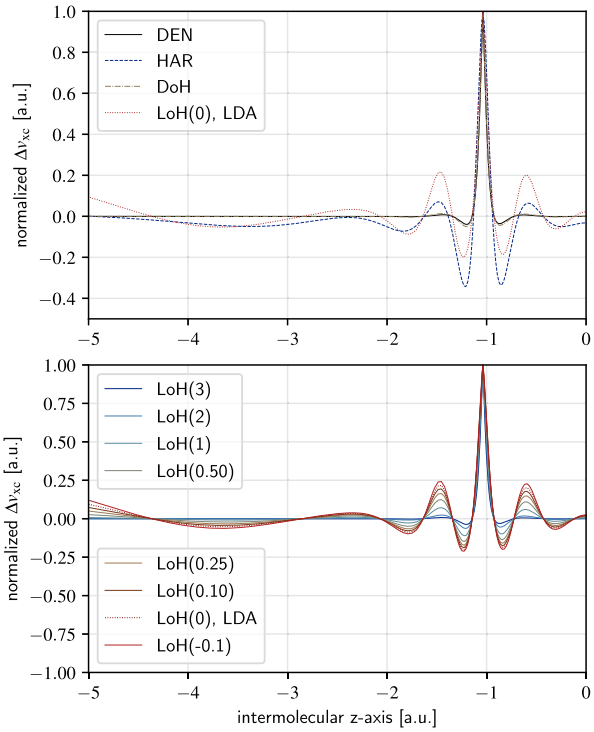

FIG. 3. Visualization of the effects of the different update rules discussed in Sec. III. The upper panel shows the update rules DEN, HAR, DoH, and LoH(0), while the lower panel shows the update rule LoH(β) for different values of β. The density difference is apparent from the DEN method. The density n' for evaluating the update rules was taken from the 200th update cycle of an inversion based on the HAR method for N_2 with the QMC-A reference density. The nuclei are at the positions $\pm 1.037\ a_0$, i.e., the plot was deliberately chosen to show one nucleus and the asymptotic region. The initial xc potential is the LDA potential from a standard self-consistent DFT calculation. The reference density is the self-consistent density of a DFT calculation with the exact exchange functional within the Krieger–Li–Iafrate (xKLI) approximation.[57] The update potentials Δv_{xc} have been normalized to their maximum value on the plotted interval for easier visual comparability.

can put any of the other above-presented update rules. For example, we denote by "HAR+DEN" an inversion in which the HAR rule was used in combination with the DEN rule, i.e., $v_{xc\sigma}^{(i+1)} = v_{xc\sigma}^{(i)} + \epsilon_{HAR}\Delta v_{xc\sigma}^{HAR} + \epsilon_{DEN}\Delta v_{xc\sigma}^{DEN}$. Due to the large values of density differences, we choose a feedback strength of $\epsilon = 0.1$ for every rule but HAR. Note that by naïve combination of update rules, their effects on the potential could cancel each other, thus slowing or even preventing convergence. However, if feedback strengths and rules are chosen carefully, the advantages of the different rules can be combined, leading to better convergence. The label HAR + LoH(0.5), e.g., corresponds to an update rule of $\Delta v_{xc\sigma} = 0.5\Delta v_{xc\sigma}^{HAR} + 0.1\Delta v_{xc\sigma}^{LoH}$ with $\beta = 0.5$. While we have defined the update rule LDA in Eq. (13), we do not report any results for this rule since it is equivalent to LoH(0) up to a constant factor.

This concludes our discussion of the update rules and their specific properties. In Sec. S2 of the supplementary material, we address further technical details of the iterative scheme. The iterative scheme for the inversion consists of two major choices: the initial xc potential and the update rule. Both affect the converged inverted potential and thus might be seen as biases or constituents of the regularization strategy. A third choice are the implementation details of the inversion procedure. While they are more likely to decide about the stability of the inversion, i.e., whether and how fast convergence is achieved, they usually have less impact on the final converged potential.

For obtaining the results that we present in Sec. VI, we chose the update rules in view of the specific task of inverting QMC densities, as explained in Sec. VI. This implies that special care was taken to ensure compatibility with statistical fluctuations in the reference density. As a general rule, a smoother reference density, e.g., one that admits to accurately take numerical derivatives, allows for more sophisticated update rules, for example, those of Refs. 19, 50, 51, and 54, while with less regular reference densities, one might have to restrain oneself to pure HAR updates.

Finally, we mention in passing that one might think about Kohn–Sham inversion from a different perspective. Instead of thinking in terms of iterative updates, one might try to exploit that even a relatively large change in v_{xc} typically still leads to a linear density response. As we have not yet tested this alternative approach in practice, and as understanding it requires a detailed analysis of the Kohn-Sham equations, we defer this topic to Secs. S.3 and S.4 in the supplementary material, which also contains Refs. 58–61.

IV. MEASURING INVERSION QUALITY

It is known from earlier work on Kohn–Sham inversion that significantly different xc potentials can map to very similar densities.[23,62] However, in the inversion process, the given reference density is the only reference point available for deciding about the quality of the inversion. Therefore, one has to find a measure of "distinctness" or "closeness" of two densities in order to define a criterion for the quality of the inverted xc potential. We here utilize two standard norms: on the one hand, the integrated absolute density difference,

$$D_1[n^{(k)}, n] = \sum_\sigma \int \left| n_\sigma^{(k)} - n_\sigma \right| \mathrm{d}^3 r \in [0, 2N], \tag{16}$$

and on the other hand, the infinity norm,

$$D_\infty[n^{(k)}, n] = \max_{\sigma,\mathbf{r}} \left| n_\sigma^{(k)}(\mathbf{r}) - n_\sigma(\mathbf{r}) \right|. \tag{17}$$

We additionally also used the Wasserstein or Kantorovich–Rubinstein metric[63,64] $W_2[n, n']$. It quantifies the minimal Euclidean transport of probability density that is necessary to align two given densities. In this way, it defines a measure of closeness beyond points-wise arguments. The measure is explicitly tailored toward comparing probability densities. We calculate the Wasserstein distance with the Python optimal transport library POT.

Although two xc potentials for a given system typically show a much larger deviation from one another than two densities, it is often not helpful to compare potentials based on the metrics D_1 and D_∞ due to the ambiguous constant offset in the potential. Even if one calculates the above-mentioned metrics with an offset minimizing the metric, the above-discussed ambiguities in the asymptotic

domains and at the nuclei limit the usefulness of comparing potentials in real space. Fortunately, the potential rather directly gives rise to electronic properties such as the eigenvalues, energies, and other observables that can be obtained from the Kohn–Sham Slater determinant. Comparing those indirect observables instead of the potentials themselves filters out features of the potential that are not relevant for the actual physical interpretation. The inspection of the Kohn–Sham eigenvalues is an elegant way to assess how different features in the xc potential affect the orbitals and their energies. For two potentials $v^{(0)}$ and $v^{(1)}$, one can compare their respective eigenvalues, $\left(\varepsilon_{\sigma i}^{(0)}\right)$ and $\left(\varepsilon_{\sigma i}^{(1)}\right)$, by the metric,

$$\Delta\varepsilon[v^{(0)}, v^{(1)}] = \frac{1}{N}\sum_{\sigma}\sum_{i=1}^{N_\sigma}\left|\left(\varepsilon_{\sigma i}^{(0)} - \varepsilon_{\sigma i-1}^{(0)}\right) - \left(\varepsilon_{\sigma i}^{(1)} - \varepsilon_{\sigma i-1}^{(1)}\right)\right|. \quad (18)$$

This metric compares the distance from one eigenvalue to the next for the two spectra and is invariant under a global potential offset. To make the metric comparable for systems with different numbers of electrons, we divide by the particle number.

As a particularly relevant observable, one could consider the HOMO eigenvalue because in exact DFT, it equals the negative of the ionization potential, $I = -\varepsilon_{\mathrm{HOMO}}$. With the eigenvalues being sensitive to a constant offset in the potential and, therefore, being connected to the asymptotic behavior of the xc potential and its high ambiguity, one cannot expect that the HOMO eigenvalue that one obtains from a potential that has been found by inversion is as such close to the ionization potential. For example, an inversion initialized by the xc potential of an LDA calculation will typically not exhibit the true $-1/|\mathbf{r}|$ behavior for $|\mathbf{r}| \to \infty$. The final inverted potential will show some mixture of decay types, stemming from the initial LDA potential in the far asymptotic domain, while in the near asymptotic domain, the update rule will have changed the potential's shape. Thus, the potential offset is hard to quantify, and therefore, the HOMO eigenvalue typically does not directly reflect the ionization potential.

However, one can define a further measure of inversion quality. Assume that the true xc potential corresponding to the reference density would be known. Then, one can exploit that for two potential-density pairs, (v_{xc}, n) and (v'_{xc}, n'), direct evaluation of J from Eq. (3) is possible, i.e., one can calculate

$$J_{\mathrm{xc}}[v_{\mathrm{xc}}, v'_{\mathrm{xc}}] = \sum_{\sigma}\int\left(v_{\mathrm{xc}\sigma} - v'_{\mathrm{xc}\sigma}\right)\left(n_\sigma - n'_\sigma\right)\mathrm{d}^3r. \quad (19)$$

Evaluating J_{xc} with an approximating potential-density pair and the true pair indicates how close the pairs approach each other. The metric scales quadratically with the distance of the pairs and reaches zero for identical pairs within the limitations discussed in Sec. II.

One cannot directly apply the metric J_{xc} to a general inversion process since the true xc potential is of course not known initially. However, one can follow Ref. 54 and "abuse" the discriminator J_{xc} by inserting the xc potentials of the last two iterations together with the current density and the reference density to define a spatially resolved convergence indicator,

$$\zeta_{\mathrm{xc}\sigma}(\mathbf{r}) = \left(v_{\mathrm{xc}\sigma}^{(k+1)}(\mathbf{r}) - v_{\mathrm{xc}\sigma}^{(k)}(\mathbf{r})\right)\left(n_\sigma(\mathbf{r}) - n_\sigma^{(k)}(\mathbf{r})\right), \quad (20)$$

with the integrated metric,

$$Z_{\mathrm{xc}} = \sum_{\sigma}\int\zeta_{\mathrm{xc}\sigma}\,\mathrm{d}^3r \quad \text{and} \quad Z_{\mathrm{xc}}^{\max} = \max_{\sigma,\,\mathbf{r}}\zeta_{\mathrm{xc}\sigma}(\mathbf{r}). \quad (21)$$

The inserted new potential and the reference density do not correspond to each other, i.e., the reference density does not stem from the potential $v_{\mathrm{xc}\sigma}^{(k+1)}$. Therefore, there are no rigorous bounds for Z_{xc} such as $J_{\mathrm{xc}} < 0$. Although Z_{xc} is a convenient number to monitor the inversion iteration, it comes with some caveats. First, a large constant potential offset from one iteration to the next might spoil the interpretation of Z_{xc}. Second, Z_{xc} can become zero even though the density difference does not vanish. This case will be encountered if the inversion cycle reaches a stationary state with $v_{\mathrm{xc}\sigma}^{(k+1)} = v_{\mathrm{xc}\sigma}^{(k)}$ due to limitations of the update rules.

V. SIMPLE INVERSE PROBLEMS

The aim of inversion techniques is to generate the Kohn–Sham potential that reproduces a given accurate reference density, and the aim of our paper here is to find the potentials that reproduce the ground-state densities that were recently obtained in quantum Monte Carlo (QMC) calculations on real-space grids.[43] It has been demonstrated that the grid approach has the advantage of allowing for a very systematic convergence—a property that is hard to achieve with the otherwise typically used localized basis sets. However, doing all-electron calculations on real-space grids is non-trivial, as the potential and the orbitals vary rapidly close to the nuclei, requiring high spatial resolution with many grid points in these regions. Therefore, the QMC calculations of Ref. 43 used a special prolate spheroidal coordinate system exploiting azimuthal symmetry, in which the points of a two-dimensional grid are densely packed close to the nuclei.[65] Consequently, our inversions must work on the same grids.

Therefore, before inverting the QMC densities, we first test how well the inversion schemes can be realized in general and on this special grid structure. Such a test can excellently be carried out by doing a quasi-trivial inversion. For example, one can do a standard DFT calculation with an approximate v_{xc}, e.g., LDA, to find the LDA ground-state density. Then, one can use the inversion schemes with this LDA density as the reference density. Ideally, the numerical inversion should then lead exactly back to the original LDA potential. By comparing the potential obtained by inversion with the original LDA potential (which is not used during the inversion process), one can directly assess the accuracy of the inverted potential by comparing the two potentials in real space and by using the measures of Eqs. (18) and (19).

A challenging system even for this type of seemingly simple inversion is the carbon atom. Its spin polarized density is determined by four spin-up electrons with the configuration $(1s2s)2p^2$ and two spin-down electrons with $(1s2s)$. The density obtained in an xKLI calculation shows a non-spherical asymptotic decay with a non-vanishing asymptotic constant along the nodal line of the HOMO orbital. As a result, the potential builds up a ridge along the nodal axis.[17,20,41,42,66,67] The LDA xc potential does not show such complicated features, and therefore, the xc potentials of LDA and xKLI are very different. With a reference xc potential (xKLI) exhibiting the correct $\propto 1/r$ asymptotic decay, an asymptotic ridge, inter-shell bumps, and quadratic behavior at the nucleus on the one hand, and an initial xc potential (LDA) with exponential decay, no

asymptotic ridge, and cusps at the nuclei on the other hand, one can set up an extraordinarily difficult inverse problem. In such a situation, one does not expect to recover the reference potential (xKLI) exactly since an asymptotic nodal ridge lies beyond the update rules' capabilities and also beyond the general limitations of the inversion mentioned in Sec. II.

With these prerequisites in mind, we inverted the xKLI density (reference density) with an initial xc potential from an LDA calculation as a "worst case test." Figure 4 shows results for the HAR, HAR+DEN, HAR+LoH(0), and HAR+LoH(1) update rules that were introduced in Sec. III. Table I lists the corresponding inversion metrics. In these calculations, the HAR update rule uses spin-resolved Hartree potentials. As expected, from the initial LDA potential, no update rule converges to the xKLI potential, but the updates lead to an intermediate potential that holds some features of the LDA and some of the xKLI potential. For example, as a consequence of how the update rules are constructed and how the initial potential was chosen, no inverted potential is quadratic at the nucleus. The inter-shell bump, where the density takes values of about 7×10^{-2}, is partially generated by all update rules. However, it has a local maximum only for HAR+LoH(0) and a saddle point for HAR+LoH(1). The nodal ridge in the target xKLI potential is not recovered by any update rule (not shown), as expected. Despite these quite noticeable differences in the real-space visualizations of the inverted xc potentials, the density differences in Table I are rather small. As a further test, we, therefore, look at the Kohn–Sham eigenvalues. Figure 5 shows that in the energy range studies here, the visible differences between the xc potentials do not influence the eigenvalues very much. A somewhat larger difference is seen only for the lowest occupied eigenvalue, which is explained by the different potential structure at the nuclear position. Otherwise, while the LDA eigenvalues differ decisively from the xKLI ones, the eigenvalues from the inverted potentials are quite similar to the xKLI eigenvalues in the examined energy range. Comparing the different inversion schemes, we find that while the HAR and HAR+DEN rules show inferior density metrics, they lead to reasonable spectral properties, with slightly worse eigenvalues for HAR. The updates from HAR+LoH(1) as well as HAR+LoH(0) further improve the inversion for the density metrics and the eigenvalue metric.

The above-mentioned methodology might be called an *inverse crime*,[23,68] i.e., evaluating the forward map with the same numerical methods that the inverse map relies on. In our case, the above-mentioned study is solely dedicated to the numerical understanding of the inverse map. In the following, we apply the inversion procedure to densities from a different numerical approach. From comparing the inversions for the carbon atom to the inversions that we did for the dimers Li_2 and N_2, we learned that the HAR method discriminating spin channels is more accurate for truly spin polarized systems, i.e., the carbon atom, than the HAR method that uses the spin-averaged potential. However, when we calculate the

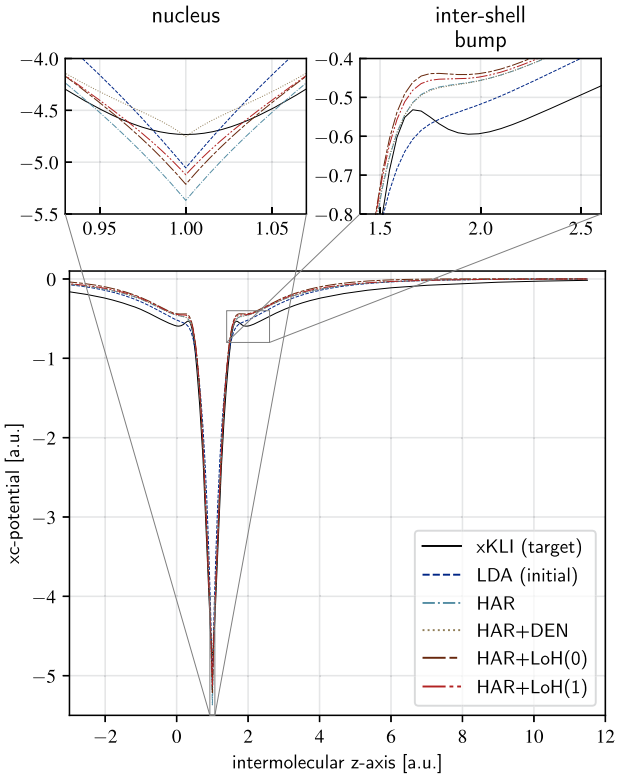

FIG. 4. "Simple" inversions for the carbon atom. The reference density originates from a standard DFT calculation using the exact exchange functional in the KLI approximation (potential given by the solid black line). The initial potential is the self-consistent LDA potential (dashed blue line). Different update rules are examined, with their xc potentials indicated by the lines of different colors and dash styles. The inverted potentials resemble only some of the features of the target xc potential. The upper insets zoom in on the nucleus and the location of an inter-shell bump.

TABLE I. Metrics of a "simple" inversion for the carbon atom: The self-consistent xKLI density was chosen as the reference density, and various inversions were performed starting from the xc potential of a self-consistent LDA calculation. The iterative improvements to the potential tend to be non-monotonic in the inversion process. Thus, we store the potential corresponding to the density that most closely resembles the reference density. The iteration in which this optimal density occurred is denoted as the optimal iteration.

Method	$D_1 \times 10^{-3}$	$D_\infty \times 10^{-2}$	$W_2 \times 10^{-5}$	$J_{xc} \times 10^{-5}$	Optimal iteration
HAR	6.28	13.4	4.52	−20.9	254
HAR+DEN	4.97	3.45	4.32	−11.7	184
HAR+LoH(0)	3.09	9.32	1.98	−6.71	373
HAR+LoH(1)	3.13	5.48	3.01	−5.43	399

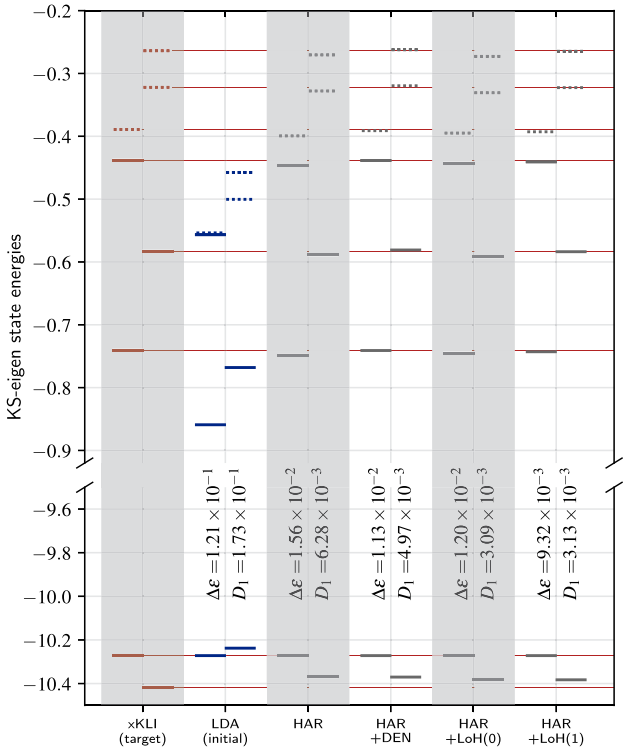

FIG. 5. Comparison of the eigenvalue spectra (Hartree units) obtained in "simple" inversions for the carbon atom. The self-consistent xKLI density was chosen as the reference density, and various inversions were performed starting from the xc potential of a self-consistent LDA calculation. Left of the respective ticks that denote each inversion method are the corresponding spin-up eigenvalues and right of it are the spin-down eigenvalues. The full lines correspond to occupied eigenvalues and the dashed lines to unoccupied ones. The eigenvalues of the xKLI calculation are marked with thin horizontal lines to ease the comparison with the eigenvalues from the inverted potentials. Note that the HOMO eigenvalues in the up-channel are doubly degenerate.

inversions for the spin unpolarized systems Li_2 and N_2 with two independent spin channels, a spin-resolved Hartree potential update decreases the accuracy slightly. Therefore, we used the spin-resolved HAR method only with truly spin polarized systems.

VI. QMC DENSITY INVERSION

The reference densities of this section stem from QMC calculations, more precisely from fixed-node diffusion Monte Carlo with a single determinant Slater–Jastrow trial wave function and the binning method to calculate the density, extrapolated to reduce the mixed estimator bias. We use the densities published in Ref. 43. As those densities are given on grids that are so large that they are infeasible for efficient inversion, we describe an interpolation scheme for reducing the size of the grids in Sec. S1 of the supplementary material.

The key challenge in inverting the QMC densities is the presence of statistical noise. This noise is point-wise in the density and, while unnoticeable to the bare eye for most regions of space, the asymptotic domain and the vicinity of the nuclei are particularly affected by noise. As discussed above for the example of an incorrect nuclear cusp, even tiny numerical artifacts might have a severe impact on the inverted potential. Therefore, we use the regularization scheme of Ref. 43 to remedy the critical statistical fluctuations of the density at the nuclei by invoking the cusp condition;[69] for details, cf. Sec. S1 of the supplementary material.

We have inverted seven different reference densities from QMC calculations. The QMC method requires starting orbitals. For the densities that are labeled "QMC-A" in the following, we chose DFT orbitals from the LDA functional in the Perdew–Wang parameterization.[70] For the densities labeled "QMC-B," we chose Kohn–Sham exact exchange orbitals with the potential in the KLI approximation. The systems under study are the nitrogen molecule (with a bond length of 2.0743 from Ref. 71), Li_2 (with a bond length of 5.051 from Ref. 71), and the carbon atom. In addition, we calculated another density based on LDA orbitals for Li_2 labeled "QMC-C." This density is entirely analog to the "QMC-A" density, but with different statistical fluctuations, i.e., a different seed for the random number generator for the QMC calculation. We carried out inversions with initial xc potentials from self-consistent LDA and xKLI calculations as well as from the Fermi–Amaldi potential. We have conducted inversions for every reference density, every initial potential, and every update rule of Sec. III with values $\beta = -0.1, 0.1, 0.25, 0.5, 1, 2,$ and 3 for the LoH update. However, only a subset of the results is reported here.

To begin with, we inspect the metrics of the inversions for the nitrogen molecule in Table II comparing different update rules and initial xc potentials.

TABLE II. Metrics of inversion for the N_2 molecule. Inversions were performed for the density QMC-B. The entry marked by an asterisk, HAR+LoH(1), is an example of a very successful inversion, as shown in Fig. 6. The rule HAR+LoH(0.5)[xKLI], marked by †, is an example of an almost instable inversion attempt. The inversion procedure did not converge with the same "soft" parameters as the other inversions; cf. Sec. S2 of the supplementary material. While for further analysis this HAR+LoH(0.5)[xKLI] inversion attempt should be discarded, it is not generally impossible to carry out successful inversions for the combination HAR+LoH(0.5)[xKLI].

Method	$v_{xc}^{(0)}$	$D_1 \times 10^{-3}$	$D_\infty \times 10^{-1}$	$W_2 \times 10^{-5}$	$Z_{xc} \times 10^{-7}$
HAR	FA	14.8	5.94	1.65	−6.2
HAR+LoH(0.0)	FA	11.1	6.38	0.81	−10.2
HAR+LoH(1.0)	FA	11.1	5.90	1.01	−7.2
HAR+DEN	FA	9.7	2.97	0.83	−5.3
HAR	LDA	16.5	8.32	1.47	−19.6
HAR+LoH(0.0)	LDA	13.8	8.10	1.15	−22.3
HAR+LoH(1.0)	LDA	12.7	7.22	1.21	−14.1
HAR+DEN	LDA	11.2	3.59	1.06	−9.2
HAR	xKLI	9.7	6.33	1.83	−1.3
HAR+LoH(−0.1)	xKLI	7.9	6.20	1.27	−2.3
HAR+LoH(0.0)	xKLI	7.6	6.11	1.24	−2.1
HAR+LoH(0.25)	xKLI	7.6	6.02	1.41	−2.0
HAR+LoH(0.5)†	xKLI	10.1	6.32	2.15	−7.6
HAR+LoH(1.0)*	xKLI	6.6	5.27	1.29	−1.5
HAR+LoH(2.0)	xKLI	7.9	4.72	1.46	−2.5
HAR+LoH(3.0)	xKLI	7.9	4.13	1.55	−2.5
HAR+DEN	xKLI	10.0	3.73	1.91	−6.1

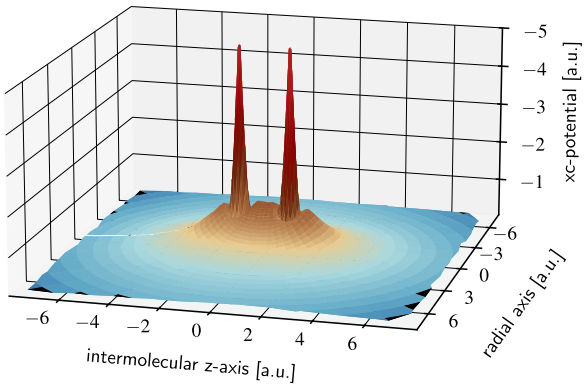

FIG. 6. Three-dimensional representation of the inverted xc potential of the nitrogen molecule in a plane through the molecular axis. The reference density QMC-B was inverted with the update rule HAR+LoH(1) starting from an initial xKLI potential. Despite the stochastic fluctuations in the reference density, the inverted potential has experienced enough regularization to be smooth. The inter-shell bumps are clearly visible around the nuclei. Note the inverted axis for the potential. The coloring was chosen to resolve the key features of the potential.

One of the nicest inversions for nitrogen is generated by HAR+LoH(1)[xKLI], i.e., an inversion using the update rule HAR+LoH(1) and an initial xc potential from the xKLI approximation. It is illustrated in Fig. 6. The inverted potential features the inter-shell bumps and the correct asymptotic decay, inherited from the xKLI potential. It is smooth and shows no artifacts even in the vicinity of the nuclei. The inverted density matches the reference density with up to $D_1/N < 0.05\%$ of the electron density correctly located. The inversion converged to $Z_{xc} < 2 \times 10^{-7}$. This qualifies this specific inversion as a success. Nevertheless, other inversions show similar metrics albeit with different features in the potential. One example is HAR+LoH(0)[xKLI], which is shown as part of Fig. 7. This inversion leads to numerical artifacts in the asymptotic domain. Another example is HAR+DEN[FA], shown as a part of Fig. 8. Here, the inversion leads to yet larger deviations and much larger fluctuations at the nuclei, but still comparable metrics.

Table II highlights some of the features of the update methods. Comparing the influence of the different starting potentials, we see that the overall best results are reached with the xKLI potential. The Fermi–Amaldi potential is still a good starting potential. The LDA shows inferior metrics. These findings are in line with straightforward expectations: xKLI is the best starting point because its potential already includes important features of the exact potential, such as proper inter-shell bumps and asymptotics (cf. Fig. 7). The Fermi–Amaldi potential misses many of these features, but still has the proper asymptotics (cf. Fig. 8). LDA is limited by its incorrect asymptotic behavior.

The differences between the different update rules are less pronounced, except for HAR and HAR+DEN. The rule HAR does not fully recover the density. We attribute this to breaking sign-consistency on a relatively large scale, imposing a strong regularization on the inversion. Its metrics reflect its limitations, while the resulting inverted potentials are smooth; see Figs. 7 and 8. With direct density feedback, the rule HAR+DEN excels to minimize the metric D_∞ compared to even the best other update rules. The

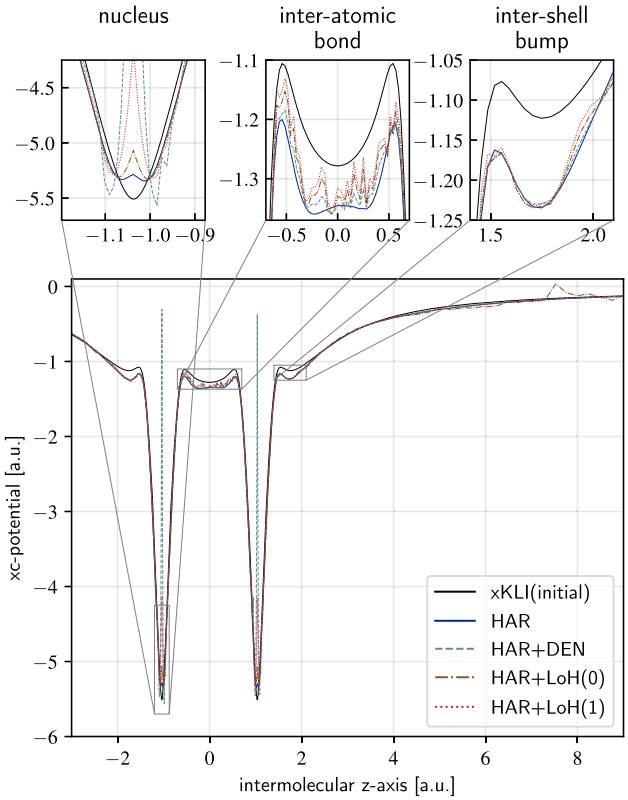

FIG. 7. Different inverted xc potentials for the density QMC-B of the N_2 molecule. Starting from the exchange-only Kohn–Sham potential in the KLI approximation (xKLI), the inverted potentials were generated by the update rules HAR, HAR+DEN, HAR+LoH(0), and HAR+LoH(1). See the main text for a detailed explanation of the different panels.

D_∞ metric is most affected by the density deviations at the nuclei, where the density takes its maximum values. Therefore, HAR+DEN best resembles the density at the nuclei. This may come at the cost of introducing artifacts in the inverted potential and lower accuracy for other properties such as eigenvalues. For the case of HAR+LoH(0.5)[xKLI], which is marked with a dagger in Table II, the iterative inversion updates almost diverged; see Fig. S5 of the supplementary material. From the metric D_∞ alone, one would hardly recognize the failure of this inversion. This observation highlights the importance of looking at different metrics to monitor the success of an inversion. The metrics D_1 and W_2 clearly mark this specific inversion as a problematic outlier.

The parameter β in the rule LoH offers a way to continuously switch between different limiting cases: On the one hand, low values of β lead to smoothness of the resulting xc potentials at the cost of slightly unresolved density differences. Higher values, on the other hand, lead to a more accurate restoration of the reference density, yet at the cost of sacrificing regularity in the xc potential and possibly introducing artifacts. For the inversions started from the xKLI potential (with the exception of the outlier), the update rule LoH(β) shows monotonously decreasing D_∞ with increasing β, bridging from the rule HAR to HAR+DEN. The different LoH rules improve

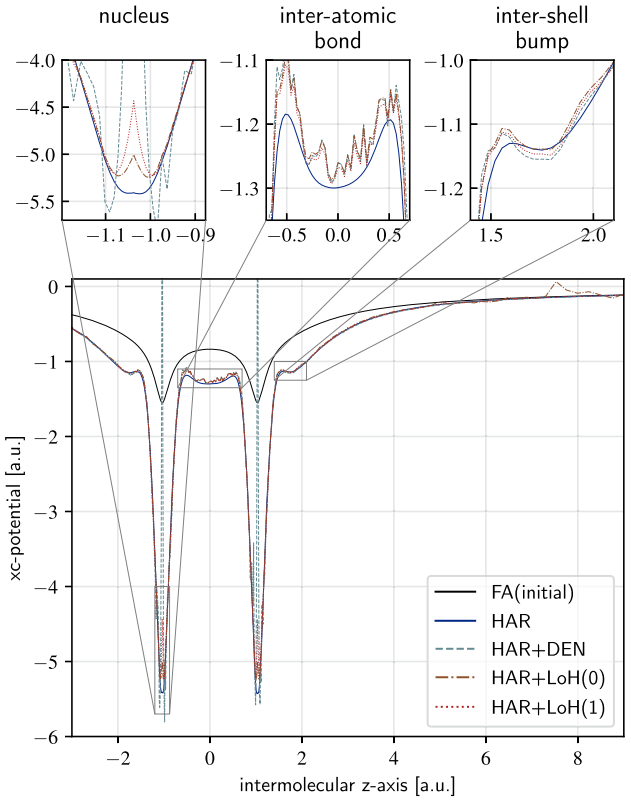

FIG. 8. Same as in Fig. 7, with the difference that here the inversions were started from the Fermi–Amaldi (FA) potential.

upon HAR and HAR+DEN in the Wasserstein metric W_2. The metric D_1 has a minimum for HAR+LoH(1) with maximum values for HAR and HAR+DEN, indicating a "sweet spot" for the parameter $\beta \approx 1$.

A visualization of some of the inversions is given in Fig. 7 for the case of starting from the Kohn–Sham exchange only potential in the KLI approximation and in Fig. 8 for the case of starting from the Fermi–Amaldi potential. Both figures have four panels to allow for a detailed comparison. The lower main panel gives an overall impression of the potentials that are reached by the different inversion calculations. The three inset panels at the top of each figure show magnifications of the three spatial regions that are particularly challenging for the inversion: the nucleus, the inter-atomic bond region, and the inter-shell bumps at the transition to the asymptotic region. The two latter regions are challenging because there the density is relatively low, yet the potential has pronounced features. The former is challenging for the reasons discussed in Sec. II.

First, let us look at the large panel showing the overall potentials. Comparing Figs. 7 and 8 reveals that the two initial potentials, xKLI and Fermi–Amaldi, each visualized by a full black line, differ substantially. Nevertheless, in both cases, the inversion procedures converge to potentials that are overall very similar. In the course of the inversion, substantial changes are made to the initial potentials. These are especially large and easy to see for the Fermi–Amaldi case: while the initial potential in Fig. 8 is relatively structureless, with the only significant features being the minima at the nuclear positions, the inverted potentials show many more features. The inter-atomic bond region has been lowered significantly, and two local maxima have appeared in it. Furthermore, at about $z \approx \pm 1.7\ a_0$, the two inter-shell bumps have appeared in the inverted potentials. Exactly the same features appear in the inverted potentials shown in Fig. 7. This demonstrates that these are genuine features of the quasi-exact v_{xc} and not remnants of the initial potentials; i.e., the iterative updates succeed in bringing out decisive features of the exact v_{xc} independent of the initial potential.

Next, let us look at the insets for the nuclear regions, i.e., the top left panels in Figs. 7 and 8. As discussed in Sec. II, the practical relevance of obtaining a good representation of v_{xc} at exactly the nuclear position is very limited because the total Kohn–Sham potential at this point is completely dominated by the diverging external potential. Nevertheless, it is interesting to see how the different update rules cover this region. The HAR+DEN update leads to a diverging potential with strong wiggles in the vicinity of the nucleus. The potentials generated by the other update rules approach the nuclear position more regularly, with the pure HAR update leading to the smoothest potential.

The superior smoothness of the HAR update is also seen when looking at the inter-atomic bond region, i.e., the top-middle panels in Figs. 7 and 8. The inverted potentials obtained with all other rules show significant wiggles in this region. If one would have done just one inversion, it would have been tempting to disregard these wiggles as numerical artifacts. However, when comparing the results from the different update rules and initial potentials, one realizes that these wiggles, although not exactly identical in the different calculations, share surprisingly similar characteristics throughout. The explanation for this observation lies in the nature of the target density: Stemming from a QMC calculation, it contains inevitable statistical noise, and the accurate representation on a real-space grid captures this noise in detail, cf. Ref. 43. Thus, what we see here is an example for the sensitivity of the Kohn–Sham inversion: when one requests that the inversion should reproduce the target density with great accuracy, the inverted potential builds up wiggles to reproduce the statistical fluctuations of the QMC density. These fluctuations also explain the slight asymmetry that is observed. As discussed previously, the HAR update is not particularly accurate and fails to reproduce the target density with high accuracy. Here, one might see this as an advantage because the HAR update, therefore, yields a smooth potential. Due to this smoothness, one can recognize the build-up of a small repulsive bump in the xc potential at $z = 0$ in Fig. 7, i.e., for the HAR[xKLI] inversion. This bump cannot be identified from any of the other inversions due to the pronounced wiggles.

Inspection of the top right panels in Figs. 7 and 8 shows that as such the outer inter-shell bump is a very stable feature appearing with a similar shape in all inversions. It is better captured by the inversions that start from the xKLI potential because the latter by itself already features an inter-shell bump. In addition, here, the statistical character of the QMC target density has some influence on the inverted potentials, but the effect is much weaker than in the inter-atomic bond region.

In order to elucidate the influence of statistical noise in the target density further, we now turn to the Li_2 molecule. For this system, we performed a series of identical inversions with two different

target densities, called QMC-A and QMC-C. Both these target densities are ground-state densities of Li_2 obtained from QMC calculations that were identical except for a different seed in the random number generator that was used to generate the (quasi) random-numbers that enter the QMC calculation. In other words, both densities were generated to represent the same physical reality, but they differ in the statistical noise that they contain. Such differences can manifest, e.g., in features such as a "bump" in the density in the far asymptotic region due to a statistical outlier in a QMC run.

The question that we want to address here is whether the statistical differences will lead to different inverted potentials. Based on the insights obtained so far, one expects the following general trends: the pure Hartree update HAR should yield a smooth potential since it averages over the fluctuations. The updates incorporating local density feedback are bound to show larger dependencies on the noise, yet will also lead to densities that are closer to the reference density.

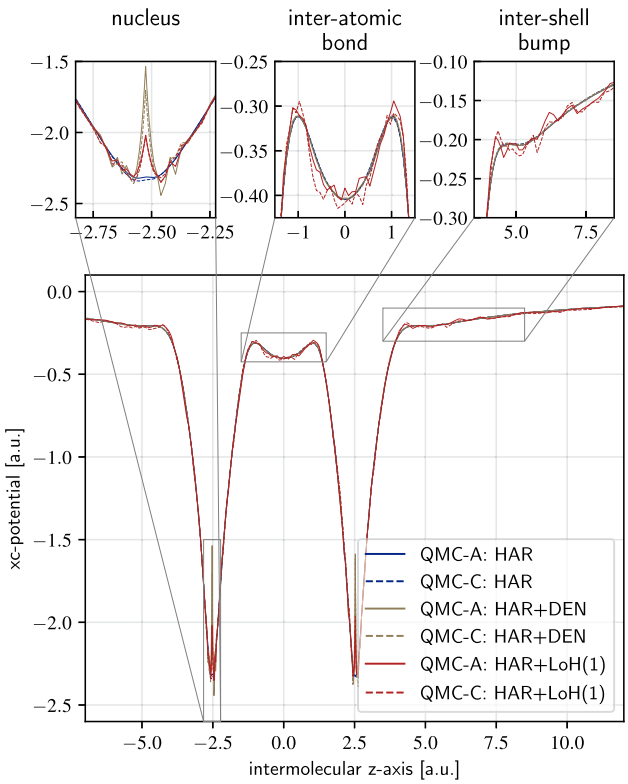

FIG. 9. Inverted xc potentials obtained for the Li_2 molecule for two different reference densities, QMC-A and QMC-C. Both were obtained from the same type of QMC calculation, and the densities differ only in their statistical fluctuations. The xc potentials generated with the update rules HAR, HAR+DEN, and HAR+LoH(1) are shown, in each case starting from the Fermi–Amaldi potential. The potentials resulting from the HAR update are not distinguishable by the eye in the main plot. The results from the HAR+DEN updates differ only marginally from the ones from the HAR update, with differences mainly visibly in the high-density regions. While the HAR+LoH(1) updates are much more "noisy," their corresponding densities are also closer to the reference density; see the metrics in Table III and Fig. 10.

In Fig. 9, we compare the inverted potentials for the rules HAR, HAR+DEN, and HAR+LoH(1) and essentially find these expectations confirmed. The potentials resulting from the HAR updates are visually indistinguishable despite the different statistical noise in the reference densities. Slight differences are visible only right at the nuclei. The rule HAR+DEN does not improve much over HAR, as the effect of the density feedback is rather weak for Li_2. The rule HAR+LoH(1) exhibits a stronger dependence on the noise pointwise, but conserves the qualitative overall features of the potential. It thus becomes clear that the ability of an update rule to locally adjust the potential via direct feedback from the density is advantageous for reference densities that are free of numerical artifacts, but can become a disadvantage in the presence of noisy reference densities. It is, however, reassuring to see that by comparing the outcome of a few different inversion calculations, one can identify which features of an inverted potential are due to statistical fluctuations in the reference density. Thus, such artificial details can be distinguished from the true characteristic features of the potential.

The just made observations and conclusions are further supported by Table III. The metrics for the corresponding inversions using the two different reference densities are close to each other, e.g., HAR+LoH(1) reaches $D_\infty = 1.96 \cdot 10^{-2}$ for QMC-A and $D_\infty = 1.88 \cdot 10^{-2}$ for QMC-C. In the Wasserstein metric W_2, there is essentially no difference between the reference densities QMC-A and QMC-C. This indicates that the numerical artifacts of the reference densities, despite having some influence on the inverted potentials, mostly influence less relevant, auxiliary characteristics of the inverted xc potentials that have little physical meaning.

One can additionally assess the consequences of the different fluctuations in the reference densities by checking how they manifest in the eigenvalues that correspond to the inverted xc potentials. For the non-trivial inversions, i.e., for the case that we have here in which the reference xc potential is not known, it is not possible to measure the quality of the eigenvalues with respect to a reference spectrum, as done previously in Fig. 5. However, one can compare the eigenvalues that are obtained from the different inverted potentials. Figure 10 shows the eigenvalues for the inverted potentials of Fig. 9 and Table III. Clearly, the eigenvalues are in overall good

TABLE III. Metrics of the inversion of two QMC densities, QMC-A and QMC-C, for the Li_2 molecule. The densities differ only in their statistical fluctuations. The numerical representations of v_{xc} obtained with update rules HAR, HAR+DEN, HAR+DoH, HAR+LoH(0), and HAR+LoH(1) are shown. In each case, the Fermi–Amaldi potential was used as the initial potential.

ref.-den.	Method	$D_1 \times 10^{-3}$	$D_\infty \times 10^{-2}$	$W_2 \times 10^{-4}$	$Z_{xc} \times 10^{-9}$
QMC-A	HAR	7.59	2.28	1.23	−4.07
QMC-A	HAR+DEN	7.12	1.73	1.21	−4.47
QMC-A	HAR+DoH	7.52	2.15	1.23	−5.32
QMC-A	HAR+LoH(0)	3.64	2.08	0.92	−3.14
QMC-A	HAR+LoH(1)	4.77	1.96	1.07	−3.49
QMC-C	HAR	8.63	2.21	1.24	−3.80
QMC-C	HAR+DEN	7.71	1.77	1.21	−3.67
QMC-C	HAR+DoH	6.58	2.03	1.17	−3.24
QMC-C	HAR+LoH(0)	2.90	1.95	0.89	−7.04
QMC-C	HAR+LoH(1)	4.26	1.88	1.06	−4.32

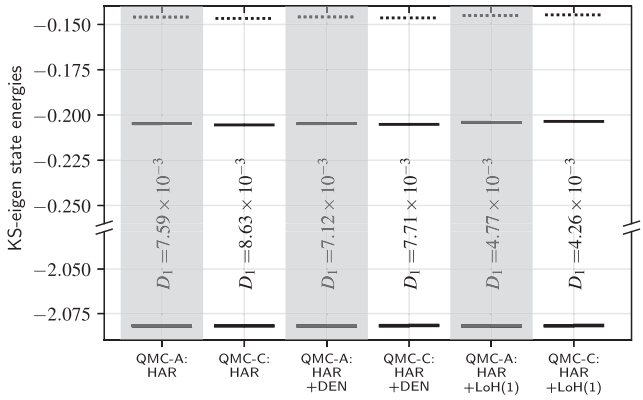

FIG. 10. Plot of the Kohn–Sham eigenvalues obtained from the potentials shown in Fig. 9. The eigenvalues have been shifted such that the lowest spin up-eigenvalues align. While originating from different xc potentials, the eigenvalues are remarkably similar. The update rules with a direct local feedback introduce a slight asymmetry in the spin channels due to picking up statistical fluctuations in the spin densities. However, despite the relatively noisy xc potentials from the HAR+LoH(1) update, the resulting eigenvalues are visually almost indistinguishable from the ones that result from the smoother xc potentials that are found with the HAR and HAR+DEN rules.

agreement. Thus, the wiggles and spikes seen in Fig. 9 hardly affect the Kohn–Sham spectrum. It is thus reassuring to see that artificial fluctuations in v_{xc} below a certain magnitude do not seem to result in noticeable physical consequences.

The finding that the physically meaningful quantities depend only relatively little on the artifacts present in the inverted xc potential can to some extent be explained by the effect of averaging over the fluctuations. Consider the reference density to be the true density, n_0, with an additional scalable artifact-density, $n = n_0 + \lambda n_a$. Then, within the linear approximation to the forward Kohn–Sham map, the corresponding true xc potential is $v_{xc0} = v_{xc} + \lambda v_{xca}$. From the Hellmann–Feynman theorem, one can conclude that the eigenvalues obey $\epsilon_i(\lambda) = \epsilon_{i0} + \lambda\langle\varphi_i|v_{xca}|\varphi_i\rangle$. Thus, the artifact contributions to the potential are averaged over the orbital's density, thus canceling out to some extent. Even without the linear approximation to the forward map, the true density response, cf. Eq. (2) of the supplementary material, also averages fluctuations of the potential in the overlap integrals $\langle\varphi_i|v_{xca}|\varphi_j\rangle$ of the response function.

Our discussion so far has focused on the statistical fluctuations that are a typical and unavoidable feature of densities that are obtained from QMC calculations. We have seen that the consequences of such fluctuations tend to be relatively mild for physically relevant quantities such as eigenvalues. However, densities can also be affected by non-statistical and systematic errors. These can result, e.g., from limited basis sets or limitations in the ansatz for the wave function such as including only a limited number of configurations in a configuration-interaction expansion. Such limitations can lead to systematic deviations that might have a more pronounced impact and might be harder to identify as the systematic errors might lead to smooth yet erroneous features in the inverted potentials.

A smooth feature that is, as extensively discussed above, difficult to accurately obtain from inversions is the asymptotic behavior of v_{xc}. For systems that have a nodal plane or line in the density of the highest occupied orbital, the asymptotic properties of v_{xc} have found particular interest: the Kohn–Sham exchange potential on the nodal surface goes to a different asymptotic constant than the potential in directions outside the nodal surface, as explicitly demonstrated both for approximations to the exact exchange potential[41] and the exact exchange potential.[42] The question of whether such features are also found in the combined exchange and correlation potential has been raised.[17] Here, we contribute to answering this question via inversions based on the QMC densities for the carbon atom from Ref. 43. Figure 11 depicts the xc potential obtained by inversion using the HAR+LoH(1) iterative update. The upper panel shows the initial xKLI potential together with the inverted potential corresponding to the QMC reference density, both along two directions: once along the nodal line of the highest occupied orbital (denoted by 0°) and once outside of the nodal surface (denoted by 90°). One can clearly see that the initial xKLI potential goes to different asymptotics in the two directions, and the inverted potential shows almost the same features. As the iterative updates are not very efficient in improving the asymptotic region, as discussed above, we show in the lower panel the same type of plots, but with the LDA

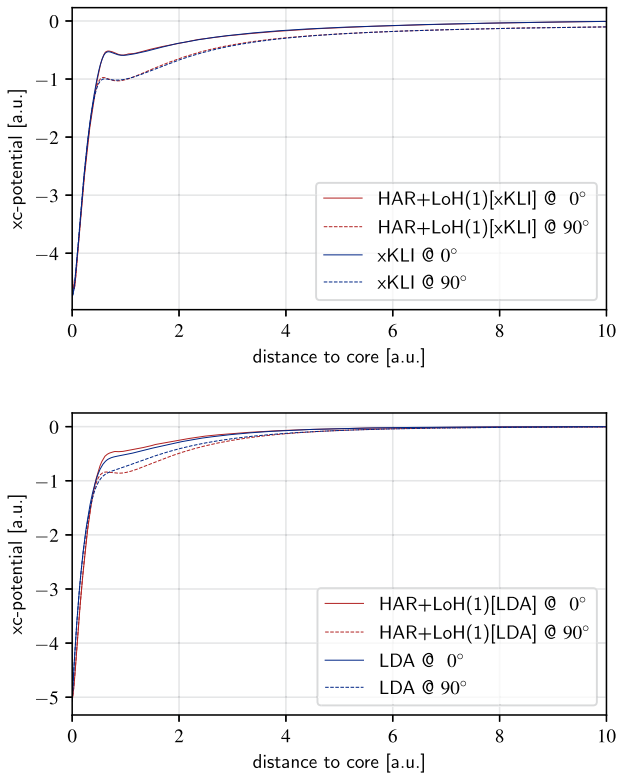

FIG. 11. xc potential obtained for the carbon atom by using the HAR+LoH(1) update rule on a reference QMC density (cf. the supplementary material). In the upper panel, the xKLI potential was used as the starting potential for the iterative updates, and in the lower panel, the LDA potential was used. The starting potentials are also depicted. All potentials are shown as a function of the distance from the nucleus for two different directions: 90° denotes a direction in which the density is dominated by the highest occupied orbital, whereas 0° denotes a direction in which the highest occupied orbital has a nodal line, i.e., vanishes.

potential as the starting point. One can see that the LDA, as expected, does not yield a faithful representation of the far asymptotic region and the iterative updates are not perfectly efficient in improving the asymptotic region. However, one can also clearly see that even when starting from the asymptotically wrong LDA, the iterative updates generate the intershell bump and introduce the same features in the region where the potential goes over into the asymptotic domain as observed in the calculations starting from the xKLI. Therefore, our calculations without doubt demonstrate that the non-vanishing asymptotic constants are not only a feature of the exchange-only potential but also prevail in the total xc potential.

As a final piece of insight into the sometimes cumbersome process of finding the xc potential corresponding to a given density, we show the evolution of the inversion over the course of the repeated application of update rules in Sec. S2 of the supplementary material.

VII. OUTLOOK: COMBINING DIFFUSION MONTE CARLO AND KOHN–SHAM INVERSION

Given the technical challenges associated with Kohn–Sham inversion, one might wonder whether it is worth the effort, i.e., what can be learnt from quasi-exact Kohn–Sham quantities? Their importance lies in a better understanding of DFT because they allow us to distinguish between features that are intrinsic to Kohn–Sham theory on the one hand, and features that are (only) consequences of certain xc functional approximations, on the other hand. One example in this respect is the Kohn–Sham eigenvalues. Although Kohn–Sham eigenvalues other than the highest occupied one do not have a rigorous physical meaning, quasi-exact eigenvalues revealed that occupied eigenvalues from the exact xc potential are very reasonable approximations to relaxed ionization energies.[7] Insight from inversion can also drive the development of improved xc approximations. An example in this respect is related to the xc derivative discontinuity: Inversion revealed that the failure of semilocal approximations to properly describe strong-field ionization is closely related to the missing xc derivative discontinuity[8] and incorporating the latter leads to significant improvements.[10,11]

In the following, we point out an additional, new direction in which inversion might play a role beyond DFT. The accuracy of DMC calculations is decisively determined by the Fermionic nodal structure of the initial trial wavefunction. The most direct approach is to use one Slater determinant as the trial wave function. The Hartree–Fock Slater determinant or a Slater determinant built from approximate Kohn–Sham, e.g., LDA, orbitals are common choices, and this choice then fixes the nodal structure. The question of whether there are functionals that yield better nodal structures is a relevant one.[72]

Thinking further, one can try to exploit that the Slater determinant built from the exact Kohn–Sham orbitals is the one unique Slater determinant that corresponds to the exact density. One might use this "density-exact" determinant as the initial trial wave function in a DMC calculation. Typically, however, the exact density and orbitals are not known *a priori*. Therefore, one can think about iteratively interlacing DMC and Kohn–Sham inversion: First, perform a Kohn–Sham calculation using some xc approximation, e.g., LDA or xKLI. Build a Slater determinant from the resulting occupied orbitals and use it as the input to a QMC calculation. The latter generates a density. Compute by inversion the Kohn–Sham potential that corresponds to this QMC density. From this potential, calculate the set of occupied Kohn–Sham orbitals that yields the QMC density. Use these orbitals to build a new Slater determinant, and use this determinant as the trial wave function in a second DMC calculation. Iterate this procedure to search for consistency between the QMC and the Kohn–Sham systems. This would yield a self-consistently determined nodal structure.

We performed a first step toward exploring this idea by going through the just described cycle once. Our test cases for this "second-order DMC scheme" were Li_2 and N_2. We first did a Kohn–Sham calculation with the xKLI potential using the previously mentioned prolate spheroidal coordinate system[43,65] with a 181×181 grid. The resulting orbitals were then used in a QMC calculation: first, a variational QMC calculation was performed to optimize the Jastrow parameter,[43] and a subsequent DMC calculation used the optimized trial wave function. The QMC density was obtained following the previously described refinement procedure. The refined density on the 181×181 grid was interpolated to a smaller gird (103×103) to allow for efficient Kohn–Sham inversion. For the inversion, we used the HAR update rule, as our main focus in these exploratory calculations was stability and smoothness of the potential. The inverted xc potential was then extrapolated to the 181×181 grid. On this grid, we again solved the Kohn–Sham equations, keeping the xc potential fixed while iterating the Hartree potential and the orbitals to self-consistency. With the resulting orbitals, a second DMC calculation was carried out, yielding the second-order DMC results.

The energy (in Hartree units) of the first DMC calculation for Li_2 using the xKLI orbitals is −14.991 43(4). The energy from the second order DMC scheme is −14.991 724(6). For N_2, the corresponding numbers are −109.516 86(3) and −109.516 77(4), respectively. We thus see that the second-order scheme leads to a slightly lower energy in the first case and a slightly higher energy in the second. However, the differences in both cases are in the range of the time step error. Therefore, from these first calculations, no clear answer emerges whether the iteration between DMC and the inverted Kohn–Sham system leads to an improvement. Furthermore, it must be kept in mind that it is a common feature of self-consistent iterations to not show monotonous improvements in the first steps. Due to the significant numerical challenges that are associated with Kohn–Sham inversion, converging the DMC-Kohn–Sham iteration is by no means trivial and convergence is not guaranteed. However, our exploratory calculations here show that first steps in this direction can be taken, and Kohn–Sham inversion might lead to insights not only into DFT but also into DMC.

VIII. SUMMARY AND CONCLUSION

In this article, we first discussed on a conceptual level the challenges that are associated with inverting the Kohn–Sham map, starting from numerically generated densities, and then looked at the advantages and limitations of specific computational realizations of the inverse Kohn–Sham map. Our particular focus was on schemes that use iterative updates to converge from an initial guess to a numerical representation of the v_{xc} that reproduces a given reference density. We introduced a new update rule that is based on density feedback with a power law weighted by the inverse Hartree potential. For the inversion of the all-electron ground-state densities that

we studied, this rule proved very successful. We investigated different iterative update rules in combination with different choices for the initial potential. Starting from a potential that incorporates the correct asymptotic behavior greatly benefits the iterative inversion. Therefore, the Fermi–Amaldi potential is a better starting point than the potential of (semi-)local approximations, such as the LDA, and the xKLI potential is a better starting point than the Fermi–Amaldi potential because it captures the inter-shell bumps and the nodal-surface features of the exact v_{xc}. To assess the quality of the results, we monitored various metrics that either give insight into density differences or highlight quantities of physical relevance, e.g., the Kohn–Sham eigenvalues.

As the Kohn–Sham inversion problem is ill-posed, one can hardly propose one specific inversion scheme that will universally work for all sorts of systems and reference densities. However, important elements for a general strategy to successfully realize the Kohn–Sham inverse map can be identified.

First, one should carefully analyze the reference density and be aware of possible limitations and artifacts. For example, when statistical noise is present, as in the case of our QMC densities, one has to expect artificial features in the inverted potential if one forces very close agreement with the reference density. Another typical limitation can be an erroneous asymptotic decay of the reference density, e.g., when it was generated by a Gaussian basis set. In such a case, one can expect possibly smooth but artificial features in the long-range part of the inverted v_{xc}.

Second, based on the character of the reference density, one should choose a set of update rules that are compatible with this character, e.g., when statistical noise is present, it is advisable to use update rules with a certain amount of regularization. As a general rule of thumb, the HAR+LoH(1) rule is a good choice for many situations.

Third, the chosen update rules should be used in combination with different initial guesses for the potential that is used to start the iterative inversion. In the choice of the initial potential, one should take into account whether particular regions are of special interest, e.g., the nuclear positions or the far asymptotics, and choose the initial guess such that it already represents these regions reasonably.

Fourth, for the different combinations of initial potentials and update rules, a series of stable inversions should be calculated and different metrics for monitoring the success of the inversions should be used.

In the course of calculating these inversions, and by comparing the results from the different calculations, one typically gains some experience, which features in the obtained v_{xc} are physical and which are artifacts. Any inverted potentials that are outliers with respect to the ensemble of the other inverted potentials should be viewed with caution. If some inverted potential shows qualitatively differing features, and the origin thereof is not explained by the construction of the update rule or the initial potential, the features might actually be artifacts.

In our experience, this strategy in practice allows one to cope with the ill-posedness of the inverse Kohn–Sham map, and following these rules, Kohn–Sham inversions can continue to contribute important insights into exact DFT, and allow us to benefit from continuing successes in the generation of accurate densities.[43,49,73] We took a first step into a possible new direction where Kohn–Sham inversion may lead to insights by exploring an iteration between DMC and the Kohn–Sham system.

SUPPLEMENTARY MATERIAL

In the supplementary material, we discuss the numerical representation of the QMC densities, we provide details about the computational realization of iterative inversion strategies, analysis of the forward Kohn–Sham map that helps gain further insight into properties of the inverse map, and a further idea for how the inverse map might be realized.

ACKNOWLEDGMENTS

We acknowledge support from the Deutsche Forschungsgemeinschaft under Grant No. 422127126 via the Bayreuth Centre for High Performance Computing (emil HPC cluster).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

A.K. and S.K. together conceptualized the work on the Kohn–Sham inversion. A.K. did all the programming, ran all calculations, and generated all tables and figures. Both authors discussed the results. A.K. wrote a first version of the manuscript and S.K. the final version. S.K. conceptualized the iteration between Diffusion Monte Carlo and Kohn–Sham inversion; A.K. did all the implementations and calculations for the second order Diffusion Monte Carlo.

Alexander Kaiser: Conceptualization (equal); Data curation (lead); Formal analysis (equal); Investigation (lead); Software (lead); Validation (lead); Visualization (lead); Writing – original draft (equal); Writing – review & editing (equal). **Stephan Kümmel**: Conceptualization (equal); Formal analysis (equal); Funding acquisition (lead); Project administration (lead); Supervision (lead); Visualization (supporting); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available within the article and its supplementary material, and additional data can be obtained from the author upon reasonable request.

REFERENCES

[1] Y. Wang and R. G. Parr, Phys. Rev. A **47**, R1591 (1993).
[2] C. J. Umrigar and X. Gonze, Phys. Rev. A **50**, 3827 (1994).
[3] Q. Zhao, R. C. Morrison, and R. G. Parr, Phys. Rev. A **50**, 2138 (1994).
[4] O. V. Gritsenko, R. van Leeuwen, and E. J. Baerends, Phys. Rev. A **52**, 1870 (1995).
[5] C. Filippi, C. J. Umrigar, and X. Gonze, J. Chem. Phys. **107**, 9994 (1997).
[6] M. J. Allen and D. J. Tozer, Mol. Phys. **100**, 433 (2002).
[7] D. P. Chong, O. V. Gritsenko, and E. J. Baerends, J. Chem. Phys. **116**, 1760 (2002).

[8]M. Lein and S. Kümmel, Phys. Rev. Lett. **94**, 143003 (2005).
[9]A. M. Teale, F. De Proft, and D. J. Tozer, J. Chem. Phys. **129**, 044110 (2008).
[10]M. Thiele, E. K. U. Gross, and S. Kümmel, Phys. Rev. Lett. **100**, 153004 (2008).
[11]A. S. d. Wijn, M. Lein, and S. Kümmel, Europhys. Lett. **84**, 43001 (2008).
[12]P. Elliott, J. I. Fuks, A. Rubio, and N. T. Maitra, Phys. Rev. Lett. **109**, 266404 (2012).
[13]J. D. Ramsden and R. W. Godby, Phys. Rev. Lett. **109**, 036402 (2012).
[14]T. Gould and J. Toulouse, Phys. Rev. A **90**, 050502 (2014).
[15]M. Thiele and S. Kümmel, Phys. Rev. Lett. **112**, 083001 (2014).
[16]O. V. Gritsenko, Ł. M. Mentel, and E. J. Baerends, J. Chem. Phys. **144**, 204114 (2016).
[17]P. Gori-Giorgi, T. Gál, and E. J. Baerends, Mol. Phys. **114**, 1086 (2016).
[18]Y. Suzuki, L. Lacombe, K. Watanabe, and N. T. Maitra, Phys. Rev. Lett. **119**, 263401 (2017).
[19]A. Kaiser and S. Kümmel, Phys. Rev. A **98**, 052505 (2018).
[20]P. Gori-Giorgi and E. J. Baerends, Eur. Phys. J. B **91**, 160 (2018).
[21]E. Fabiano, S. Śmiga, S. Giarrusso, K. J. Daas, F. Della Sala, I. Grabowski, and P. Gori-Giorgi, J. Chem. Theory Comput. **15**, 1006 (2019).
[22]R. Garrick, A. Natan, T. Gould, and L. Kronik, Phys. Rev. X **10**, 021040 (2020).
[23]Y. Shi and A. Wasserman, J. Phys. Chem. Lett. **12**, 5308 (2021).
[24]J. Erhard, E. Trushin, and A. Görling, J. Chem. Phys. **156**, 204124 (2022).
[25]Y. Shi, A. Chavez, and H. V. Wasserman, WIREs Comput Mol Sci. **12**, e1617 (2022).
[26]V. Khanna, B. Kanungo, J. Hatch, J. Kammeraad, and P. M. Zimmerman, J. Phys. Chem. A **129**, 4162–4173 (2025).
[27]S. Fauser, E. Trushin, and A. Görling, J. Chem. Phys. **162**, 164108 (2025).
[28]J. Kirkpatrick, B. McMorrow, D. H. P. Turban, A. L. Gaunt, J. S. Spencer, A. G. D. G. Matthews, A. Obika, L. Thiry, M. Fortunato, D. Pfau, L. R. Castellanos, S. Petersen, A. W. R. Nelson, P. Kohli, P. Mori-Sánchez, D. Hassabis, and A. J. Cohen, Science **374**, 1385 (2021).
[29]R. Nagai, R. Akashi, and O. Sugino, npj Comput. Mater. **6**, 43 (2020).
[30]P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964).
[31]A. Kirsch, *An Introduction to the Mathematical Theory of Inverse Problems*, *Applied Mathematical Sciences* (Springer, New York, NY, 2011), Vol. 120.
[32]J. Hadamard, Princeton Univ. Bull. **13**, 49 (1902).
[33]A. K. Theophilou, J. Phys. C: Solid State Phys. **12**, 5419 (1979).
[34]O. V. Gritsenko and E. J. Baerends, J. Chem. Phys. **120**, 8364 (2004).
[35]P. Strange, *Relativistic Quantum Mechanics: With Applications in Condensed Matter and Atomic Physics* (Cambridge University Press, 1998).
[36]Q. Zhao and R. G. Parr, Phys. Rev. A **46**, 2337 (1992).
[37]T. W. Hollins, S. J. Clark, K. Refson, and N. I. Gidopoulos, J. Phys.: Condens. Matter **29**, 04LT01 (2017).
[38]Q. Wu and W. Yang, J. Chem. Phys. **118**, 2498 (2003).
[39]C.-O. Almbladh and U. von Barth, Phys. Rev. B **31**, 3231 (1985).
[40]M. Levy, J. P. Perdew, and V. Sahni, Phys. Rev. A **30**, 2745 (1984).
[41]F. Della Sala and A. Görling, Phys. Rev. Lett. **89**, 033003 (2002).
[42]S. Kümmel and J. P. Perdew, Phys. Rev. B **68**, 035103 (2003).
[43]A. Kaiser and S. Kümmel, J. Chem. Phys. **162**, 134108 (2025).
[44]M.-C. Kim, E. Sim, and K. Burke, Phys. Rev. Lett. **111**, 073003 (2013).
[45]T. Aschebrock, T. Lebeda, M. Brütting, R. Richter, I. Schelter, and S. Kümmel, J. Chem. Phys. **159**, 234107 (2023).
[46]T. Lebeda, T. Aschebrock, and S. Kümmel, Phys. Rev. Lett. **133**, 136402 (2024).
[47]S. Fürst and M. Kaupp, J. Chem. Theory Comput. **19**, 3146 (2023).
[48]M. Brütting, H. Bahmann, and S. Kümmel, J. Chem. Phys. Rapid Commun. **160**, 181101 (2024).
[49]L. Cheng, P. B. Szabó, Z. Schätzle, D. P. Kooi, J. Köhler, K. J. H. Giesbertz, F. Noé, J. Hermann, P. Gori-Giorgi, and A. Foster, J. Chem. Phys. **162**, 034120 (2025).
[50]I. G. Ryabinkin, S. V. Kohut, and V. N. Staroverov, Phys. Rev. Lett. **115**, 083001 (2015).
[51]E. Ospadov, I. G. Ryabinkin, and V. N. Staroverov, J. Chem. Phys. **146**, 084103 (2017).
[52]T. Gould, "Toward routine Kohn-Sham inversion using the "Lieb-response" approach," J. Chem. Phys. **158**, 064102 (2023).
[53]K. Peirs, D. Van Neck, and M. Waroquier, Phys. Rev. A **67**, 012505 (2003).
[54]A. Kumar, R. Singh, and M. K. Harbola, J. Phys. B: At., Mol. Opt. Phys. **52**, 075007 (2019).
[55]L. Boccardo and T. Gallouët, J. Funct. Anal. **87**, 149 (1989).
[56]M. G. Medvedev, I. S. Bushmarinov, J. Sun, J. P. Perdew, and K. A. Lyssenko, Science **355**, 49 (2017).
[57]J. B. Krieger, Y. Li, and G. J. Iafrate, Phys. Rev. A **45**, 101 (1992).
[58]A. D. Becke, J. Chem. Phys. **98**, 5648 (1993).
[59]T. Körzdörfer and S. Kümmel, Phys. Rev. B **82**, 155206 (2010).
[60]N. Sai, P. F. Barbara, and K. Leung, Phys. Rev. Lett. **106**, 226403 (2011).
[61]V. Atalla, M. Yoon, F. Caruso, P. Rinke, and M. Scheffler, Phys. Rev. B **88**, 165122 (2013).
[62]A. Aouina, M. Gatti, S. Chen, S. Zhang, and L. Reining, Phys. Rev. B **107**, 195123 (2023).
[63]L. V. Kantorovich, Manag. Sci. **6**, 366 (1960).
[64]L. V. Kantorovich, J. Math. Sci. **133**, 1381 (2006).
[65]A. Makmal, S. Kümmel, and L. Kronik, J. Chem. Theory Comput. **5**, 1731 (2009).
[66]T. Schmidt, E. Kraisler, L. Kronik, and S. Kümmel, Phys. Chem. Chem. Phys. **16**, 14357 (2014).
[67]T. Aschebrock, R. Armiento, and S. Kümmel, Phys. Rev. B **95**, 245118 (2017).
[68]D. Colton and R. Kress, *Inverse Acoustic and Electromagnetic Scattering Theory*, *Applied Mathematical Sciences* (Springer International Publishing, Cham, 2019), Vol. 93.
[69]T. Kato, Commun. Pure Appl. Math. **10**, 151 (1957).
[70]J. P. Perdew and Y. Wang, Phys. Rev. B **45**, 13244 (1992).
[71]K. P. Huber and G. Herzberg, in *Molecular Spectra and Molecular Structure: IV. Constants of Diatomic Molecules*, edited by K. P. Huber and G. Herzberg (Springer, Boston, MA, 1979), pp. 8–689.
[72]F. A. Reboredo and P. R. C. Kent, Phys. Rev. B **77**, 245110 (2008).
[73]E. Trushin, S. Fauser, A. Mölkner, J. Erhard, and A. Görling, Phys. Rev. Lett. **134**, 016402 (2025).