Eur. Phys. J. B (2018) 91: 235
https://doi.org/10.1140/epjb/e2018-90276-4

THE EUROPEAN PHYSICAL JOURNAL B

Regular Article

Numerical construction of the density-potential mapping*

Soeren E.B. Nielsen[1,a], Michael Ruggenthaler[1], and Robert van Leeuwen[2]

[1] Max Planck Institute for the Structure and Dynamics of Matter and Center for Free-Electron Laser Science, Luruper Chaussee 149, 22761 Hamburg, Germany
[2] Department of Physics, Nanoscience Center, University of Jyväskylä, 40014 Jyväskylä, Finland

Received 23 April 2018 / Received in final form 7 July 2018
Published online 8 October 2018
© The Author(s) 2018. This article is published with open access at Springerlink.com

Abstract. We demonstrate how a recently developed method Nielsen et al. [Nielsen et al., EPL **101**, 33001 (2013)] allows for a comprehensive investigation of time-dependent density functionals in general, and of the exact time-dependent exchange-correlation potential in particular, by presenting the first exact results for two- and three-dimensional multi-electron systems. This method is an explicit realization of the Runge–Gross correspondence, which maps time-dependent densities to their respective potentials, and allows for the exact construction of desired density functionals. We present in detail the numerical requirements that makes this method efficient, stable and precise even for large and rapid density changes, irrespective of the initial state and two-body interaction. This includes among others the proper treatment of low density regions, a subtle interplay between numerical time-propagation and zero boundary conditions, the choice of time-stepping strategy, and an error damping mechanism based on both the density and current density residuals. These considerations are also relevant for computing time-independent density-functionals and lead to a more precise implementation of quantum mechanics in general, which is mainly relevant for cases in which there is notable contact with a boundary or when the low density regions matter.

1 Introduction

Computer simulations of physics at the atomic scale have become an increasingly important and integral part in the process of inventing, designing and optimizing many modern devices. The appeal is obvious, as ideally, one could for example examine all the processes that take place in a computer chip or when a drug is administered by just running a computer program, thereby saving the need to perform slow and expensive laboratory tests. In practice, such simulations are however strongly limited by the fact that their wave function formulations quickly become computationally intractable with increasing system size and significant approximations are therefore needed to overcome this bottleneck, and to allow for a study of realistic systems. Presently most simulations thus rely on density-functional theory (DFT) [1] and time-dependent density-functional theory (TDDFT) [2,3] since these methods keep an unusually high level of precision while at the same time reducing the computational cost from scaling exponentially to linearly with system size. Even so, they still reduce the precision significantly as they, especially in the time-dependent case, rely on rather simplistic approximations to the so-called Hartree exchange-correlation (Hxc) potential (or simply the exchange-correlation potential if we subtract the Hartree part). Much research has therefore already gone into this quantity [4–9], but in particular in the time-dependent case further investigations are still needed to achieve chemical accuracy and to extend the range of simulations. Adequate tools to perform such investigations are therefore essential and in this work we present a numerical method that allows to study the exact Hxc potential in the time-dependent case.

In a recent work [10] we presented a numerical inversion method to construct the external potential that, for a given initial state, produces a prescribed time-dependent density in an interacting quantum many-body system. This technique is the first general method of its kind [2] and only few other methods exist [11–15]. It combines ideas from tracking [16,17] (a local control theory also known as inverse control) with a fixed point formulation of TDDFT [18–20] and forms a highly efficient method which is precise and stable even for large and rapid density variations regardless of the initial state and two-electron interaction. It furthermore allows for the computation of the many-body wave function and consequently any property as a functional of the initial state and density, including the exact Hxc potential. This is an explicit realization of the Runge–Gross result [21–23] that any observable is a functional of the density and initial state, which forms the foundation of TDDFT. Our numerical

* Contribution to the Topical Issue "Special issue in honor of Hardy Gross", edited by C.A. Ullrich, F.M.S. Nogueira, A. Rubio, and M.A.L. Marques.

[a] e-mail: soerenersbak@hotmail.com

implementation thereof thus allows us to investigate virtually all fundamental aspects of TDDFT in a unique way. This includes the exact benchmarking of present approximations to the Hxc potential, the study of exact features of the Hxc potential for simple systems, and the study of how more exotic functionals, like the kinetic energy, depend on the density for use in orbital free TDDFT. In time, this will hopefully lead to better, more physical, approximations to the Hxc potential, and thereby an even wider use of simulations to complement laboratory work.

Inversions can be rather intricate to perform and therefore we will in this work focus on some key ingredients needed for a successful implementation. This work therefore provides significantly more details than presented in our earlier work [10]. This includes a discussion of the subtle interplay between numerical time-propagation [24] and the use of zero boundary conditions on the wave function, and an in-depth analysis of why and how one must treat the low density regions carefully. It also comprises an analysis of why we deliberately used an implicit time-stepping procedure in [10] over an explicit as proposed in [14], and a discussion on how to get a crucial error damping mechanism by using both the density and current residuals, as well as many other details. Our discussion of the low density regions is also very relevant to time-independent inversions like the ones discussed in [25,26]. At the same time we also explicitly extend the method to arbitrary dimensions such that an in-depth study of the Hxc potential for realistic, i.e., two and three dimensional systems, becomes possible. We highlight this capability by showing first results for such higher-dimensional systems.

Finally, instead of using tracking to find the potential that determines a desired density-path, like our method does, an interesting alternative is to use optimal control theory [15]. This provides some further options to penalise unwanted features in the potential. However, as tracking generally suffices for the task at hand at a significantly lower numerical cost it is generally the choice to prefer. We therefore stay with this choice throughout this article.

Outline - In Section 2 we first provide a brief overview of TDDFT in the Kohn–Sham (KS) formulation to set the background. We then explain a stripped down version of our inversion algorithm in order to demonstrate that the basic logic is rather simple and nothing but a variation of tracking. Readers only interested in the basic ideas may right after this section skip to the examples Section 7.

We stress, however, that such a naive implementation of the method is very unstable. This relates not only to the method itself, but also to an additional necessity to implement quantum mechanics more precisely than one usually does. We therefore discuss these aspects first in detail before we present the final method. Some of these aspects are also relevant for users, not just developers. For example, one cannot just use a ground state from a code that does not adhere to this rigour as an initial state in an inversion.

In Section 3 we briefly present the main ideas behind a typical implementation of the time-dependent Schrödinger equation (TDSE) to establish the numerical framework. We further point out that one must be careful how to implement zero boundary conditions for time-propagation, as it is not enough to require that the wave function is zero at the boundary if one uses analytic functions (as discussed in much more detail in Appendix A). In Section 4 we demonstrate why it is important to also treat the low density regions accurately (and in Appendix B we show how to achieve this precision even in extreme cases). With these things in place, we can then finally address the inversions, and the good news is that if one is cautious to follow the guidelines mentioned above, it becomes relatively easy to perform the inversions even in extreme cases. We present the main aspects of our inversion methods in Sections 5 and 6 for respectively time-independent and time-dependent inversions (time-independent inversions compute the potential that has a specific density as ground state density). In Section 7 we then show results for 2D and 3D systems, as well as a series of results for 1D systems, before we conclude in Section 8.

2 Overview

Throughout this paper we study the class of N-electron systems governed by the time-dependent non-relativistic spin-free Hamiltonian (atomic units are used throughout)

$$\hat{H}(t)=\hat{T}+\hat{W}+\hat{V}(t)=-\frac{1}{2}\sum_{i=1}^{N}\nabla_i^2+\sum_{i>j=1}^{N}w(\mathbf{r}_{ij})+\sum_{i=1}^{N}v(\mathbf{r}_i t), \tag{1}$$

where $\hat{T}$ is the kinetic energy operator, $\hat{W}$ the two-electron interaction operator, $\hat{V}(t)$ the external potential operator, ∇_i^2 the Laplacian with respect to the spatial coordinate $\mathbf{r}_i$ and $\mathbf{r}_{ij}=\mathbf{r}_i-\mathbf{r}_j$. We further denote the initial state of such a system by $|\Psi_0\rangle=|\Psi(t_0)\rangle$, and the state $|\Psi(t)\rangle$ then evolves according to the TDSE,

$$i\partial_t|\Psi(t)\rangle=\hat{H}(t)|\Psi(t)\rangle\,. \tag{2}$$

Given the state $|\Psi(t)\rangle$ we can then compute any desired property by evaluating the expectation value of the operator of interest. For example, the density $n(\mathbf{r}t)$ and current density $\mathbf{j}(\mathbf{r}t)$ are the expectation values of the density and current density operators

$$\hat{n}(\mathbf{r})=\sum_{i=1}^{N}\delta(\mathbf{r}-\mathbf{r}_i), \tag{3}$$

$$\hat{\mathbf{j}}(\mathbf{r})=\frac{1}{2\mathrm{i}}\sum_{i=1}^{N}\left(\delta(\mathbf{r}-\mathbf{r}_i)\overrightarrow{\nabla}_i-\overleftarrow{\nabla}_i\delta(\mathbf{r}-\mathbf{r}_i)\right). \tag{4}$$

Numerically, the computational cost of solving the exact TDSE scales exponentially with the system size, limiting us to systems with only a few degrees of freedom.

In contrast, for a non-interacting system (i.e., $\hat{W}=0$) the computational cost scales only linearly with system size assuming the initial state $|\Phi_0\rangle$ is a Slater determinant (or a linear combination of a few Slater determinants), since the time-propagated state $|\Phi(t)\rangle$ remains a Slater determinant (or a linear combination thereof) and we can therefore represent it by one-particle orbitals $\phi_i(\mathbf{r}t)$. In

this case we therefore only need to propagate each of these orbitals separately using a single-particle Hamiltonian

$$\hat{h}_s(t) = -\tfrac{1}{2}\nabla^2 + v_s(\mathbf{r}t),$$

where we employed the usual convention to denote a potential that acts on a non-interacting system by $v_s(\mathbf{r}t)$. All properties of the system can then also be expressed in terms of the orbitals $\phi_i(\mathbf{r}t)$. The density is for example (typically) given by

$$n(\mathbf{r}t) = \sum_{i=1}^{N} |\phi_i(\mathbf{r}t)|^2.$$

The cost of computing one-body properties thus also scales linearly with the particle number such that we can treat very large non-interacting systems if we limit us to such properties. The computation of many-body properties is still more expensive, like quadratic for the pair density, but still computationally cheaper than for interacting systems.

2.1 Runge–Gross mapping and Kohn–Sham theory

The computational advantages of non-interacting systems are exploited in the KS approach to TDDFT [2,27]. In this approach a non-interacting system is introduced with the same density as that of the interacting system of interest. To define this system properly we first define the Runge–Gross mapping for general interacting systems.

The time-propagation of an interacting system with external potential $v(\mathbf{r}t)$ and initial state $|\Psi_0\rangle$ produces a time-dependent many-body wave function as a functional of the potential and the initial state. A subsequent evaluation of the density expectation value produces a mapping $(\Psi_0, v) \to n$ from potentials $v(\mathbf{r}t)$ and initial states $|\Psi_0\rangle$ to densities $n(\mathbf{r}t)$ which we denote by $n[\Psi_0, v]$. The Runge–Gross mapping $(\Psi_0, n) \to v$ is then defined as the inverse of the mapping $n[\Psi_0, v]$ and can be shown to exist under mild conditions on the potentials and initial states [28]. This implies that for a fixed initial state $|\Psi_0\rangle$ and two-body interaction there is a unique potential $v[\Psi_0, n]$ (up to a purely time-dependent gauge function $c(t)$) that yields a prescribed density $n(\mathbf{r}t)$ by solution of the TDSE. The initial state is required to yield the prescribed initial density $n(\mathbf{r}t_0)$ and time-derivative of the density $\partial_t n(\mathbf{r}t_0)$ (which can be calculated by evaluating the expectation value of the current density with respect to the initial state since $\partial_t n(\mathbf{r}t_0) = -\nabla \cdot \mathbf{j}(\mathbf{r}t_0)$).

The Runge–Gross mapping exists for systems with a large class of two-body interactions. If we now consider the special case of a system in which the two-body interaction is identically zero we deduce the existence of a potential $v_s[\Phi_0, n]$ that yields the density $n(\mathbf{r}t)$ for an initial state $|\Phi_0\rangle$ in a non-interacting system. The functionals $v[\Psi_0, n]$ and $v_s[\Phi_0, n]$ can be used to define the so-called Hxc potential by

$$v_{\mathrm{Hxc}}[\Psi_0, \Phi_0, n] = v_s[\Phi_0, n] - v[\Psi_0, n] \tag{5}$$

(often further the exchange-correlation potential $v_{\mathrm{xc}}(\mathbf{r}t)$ is defined by subtracting the Hartree potential from $v_{\mathrm{Hxc}}(\mathbf{r}t)$ but this is not relevant for the present work). This functional plays an important role in the KS construction in which we can calculate the density of an interacting system with prescribed external potential $v_{\mathrm{ext}}(\mathbf{r}t)$ by time-propagation of a non-interacting system with initial state $|\Phi_0\rangle$ and potential

$$v_{\mathrm{KS}}[\Psi_0, \Phi_0, n, v_{\mathrm{ext}}] = v_{\mathrm{ext}} + v_{\mathrm{Hxc}}[\Psi_0, \Phi_0, n],$$

which is called the KS potential. To see that this procedure indeed yields the desired density we note that self-consistency is achieved, by definition, for the density $n_{\mathrm{sc}}(\mathbf{r}t)$ that satisfies

$$v_{\mathrm{KS}}[\Psi_0, \Phi_0, n_{\mathrm{sc}}, v_{\mathrm{ext}}] = v_s[\Phi_0, n_{\mathrm{sc}}].$$

From the definition of the functional v_{Hxc} we deduce that this is exactly true whenever $v_{\mathrm{ext}} = v[\Psi_0, n_{\mathrm{sc}}]$ or, equivalently, when $n_{\mathrm{sc}} = n[\Psi_0, v_{\mathrm{ext}}]$.

In practical applications of the KS formalism we need an approximation for $v_{\mathrm{Hxc}}[\Psi_0, \Phi_0, n]$. The dependence of this functional on the initial states is only known in very specific cases [6] and is generally neglected. In the special, but important, case that the initial states $|\Psi_0\rangle$ and $|\Phi_0\rangle$ are ground states they are functionals themselves of the ground state density and then v_{Hxc} becomes a functional of the density only. However, even in this case approximations for v_{Hxc} are usually only available in the adiabatic approximation and generally taken to depend on the instantaneous density or orbitals while ignoring the dependency on the density at earlier times. In fact, one mostly uses approximations designed for ground state DFT in TDDFT in an adiabatic approximation. The precision of such approximations is of course limited, but still qualitative agreement is at least usually obtained for a restricted time interval (dependent on how rapid the dynamics is). There are, however, many situations [2,7,10] in which it is necessary to go beyond the adiabatic approximation. In the design of such approximations the ability to compare to an exact benchmark is of great value and therefore in this work we present a numerical method to explicitly construct exact time-dependent density functionals.

In the remainder of this work we show how, for a prescribed time-dependent density $n(\mathbf{r}t)$ and initial state $|\Psi_0\rangle$, we can numerically construct the potential $v[\Psi_0, n]$ for any interacting system. The non-interacting potential $v_s[\Phi_0, n]$ can also be obtained since it corresponds to a special case of the same method for a vanishing two-body interaction. A quantity of special interest is the Hxc potential which can be obtained from its definition in equation (5). Another common situation is that we have solved the TDSE of an interacting problem with a given external potential v and initial state Ψ_0 and we want to study the corresponding KS system to develop and benchmark functionals in TDDFT. In this case the interacting density $n[\Psi_0, v]$ is known and the corresponding Hxc potential is obtained from $v_{\mathrm{Hxc}}[\Psi_0, \Phi_0, n[\Psi_0, v]] = v_s[\Phi_0, n[\Psi_0, v]] - v$, since this is the

Hxc potential that yields the density $n[\Psi_0, v]$ in an exact KS propagation. The wave functions $|\Psi[\Psi_0, n]\rangle$ respectively $|\Phi[\Phi_0, n]\rangle$ are also obtained as our inversion method steps through time, but are usually not stored (they can always be recomputed later by a propagation with $v[\Psi_0, n]$ respectively $v_s[\Phi_0, n]$). This allows us to also compute a large class of other density functionals as well, like the kinetic energy or different contributions to the Hxc potential, by the evaluation of expectation values of operators of interest using either the interacting or non-interacting wave function.

2.2 Basic inversion method

In this section we present an outline of the main ideas behind our inversion method while the numerical details to make it highly efficient, stable and precise will follow in later sections. The starting point of the discussion is the quantum fluid reformulation of the TDSE, which consists of the equations of motion of the density and the current density

$$\partial_t n(\mathbf{r}t) = -\nabla \cdot \mathbf{j}(\mathbf{r}t), \tag{6}$$
$$\partial_t \mathbf{j}(\mathbf{r}t) = -n(\mathbf{r}t)\nabla v(\mathbf{r}t) - \mathbf{Q}(\mathbf{r}t). \tag{7}$$

Here $\mathbf{Q}(\mathbf{r}t) = \langle\Psi(t)|\hat{\mathbf{Q}}(\mathbf{r})|\Psi(t)\rangle$ is the expectation value of the internal local force operator $\hat{\mathbf{Q}}(\mathbf{r}) = \mathrm{i}[\hat{\mathbf{j}}(\mathbf{r}), \hat{T} + \hat{W}]$ with respect to the many-body wave function $|\Psi(t)\rangle$ that solves the TDSE with potential $v(\mathbf{r}t)$. By the Runge–Gross mapping, $\mathbf{Q}[\Psi_0, n](\mathbf{r}t)$ is a functional of the density and the initial state and has the physical interpretation of an internal force density. This furnishes an exact density-functional reformulation of the quantum many-body problem in terms of a quantum fluid system. A good approximation for the functional $\mathbf{Q}[\Psi_0, n](\mathbf{r}t)$ would allow for the calculation of the density and current density without the necessity of solving the TDSE. We use equations (6) and (7) in a different way instead. Combining the two equations leads to the expression

$$-\nabla \cdot (n(\mathbf{r}t)\nabla v(\mathbf{r}t)) = q(\mathbf{r}t) - \partial_t^2 n(\mathbf{r}t), \tag{8}$$

where $q(\mathbf{r}t) = \nabla \cdot \mathbf{Q}(\mathbf{r}t)$. This equation is central to our inversion method as it directly connects the density and potential. The idea is to insert the density of interest into this equation and then use it to solve for the potential in which we regard $q[\Psi_0, v]$ as a functional of the potential $v(\mathbf{r}t)$. In earlier work [18] we showed that the potential obtained in this way indeed yields the prescribed density when we use it in the TDSE. The basic strategy of our inversion method to determine this potential is to use a time-stepping approach from local control theory known as tracking [16].

To keep the discussion as simple as possible, let us consider the Euler method of time-stepping on an equidistant time-grid with $t_n = n\Delta t$ where n is an integer. We start by computing $v(\mathbf{r}t_0)$ from the initial state $|\Psi(t_0)\rangle$ and density $n(\mathbf{r}t)$ that we prescribed. If we compute $q(\mathbf{r}t_0)$ from $|\Psi(t_0)\rangle$ then at t_0 both terms on the right hand side of equation (8) are known (since the density is known at all times) and we can solve the equation for $v(\mathbf{r}t_0)$ for which there exist efficient solvers (see Sect. 6.2). We then use Euler time-stepping with the corresponding Hamiltonian $\hat{H}(t_0)$ to obtain $|\Psi(t_1)\rangle$ as $|\Psi(t_1)\rangle = |\Psi(t_0)\rangle - \mathrm{i}\Delta t\hat{H}(t_0)|\Psi(t_0)\rangle$. Using $|\Psi(t_1)\rangle$ we can repeat the whole procedure, i.e., we calculate $q(\mathbf{r}t_1)$ and subsequently $v(\mathbf{r}t_1)$ and $\hat{H}(t_1)$, which is used to perform a new time step. Continuing this way we step through all the time points t_n thereby obtaining the desired potential $v(\mathbf{r}t_n)$ on the specified time grid. In summary: Given $|\Psi(t_{n-1})\rangle$ we compute $v(\mathbf{r}t_{n-1})$ and use it to step to $|\Psi(t_n)\rangle$. Since we need to make each time step only once the method is simple and very efficient.

The more sophisticated method that we will present later uses a similar strategy that is only modified slightly to prevent errors from building up over time. It also uses a more advanced time-stepping method than the Euler method making it only a few times more expensive than a normal time-propagation when the potential is known (instead of the density). It has the further advantage that it avoids the direct use of the complicated quantity $q(\mathbf{r}t)$ and instead employs quantities that can be expressed in terms of $v(\mathbf{r}t)$ and $n(\mathbf{r}t)$ (and possible $\mathbf{j}(\mathbf{r}t)$ for stability). This also means that our method can be combined with any spatial representation or propagation method, as long as we can propagate the density (though for approximate methods the inversion procedure may become unstable). Especially, for non-interacting systems, we can usually work in terms of orbitals, and the method is much more efficient.

The basic inversion procedure to calculate the potential has an appealing simplicity, but it is important to stress that a naive implementation of the method is highly unstable even for moderately large or rapid density changes. While the Euler method is mathematically guaranteed to converge for small enough time-steps, it usually fails in practice within the inversion method due to discretization and round-off errors.[1] One of the causes of the instability is that the density at a given spatial point may change in value by orders of magnitude. We therefore have to be very cautious to preserve a large numerical accuracy in time. For example, the density at some point in space could change from a value of 10^{-2} to a value of 10^{-10} at a later time. Consequently a small error of, for instance, 10^{-8} in the density which is regarded as small initially and which is preserved in time may be regarded as a large error at a later time when it can become much larger than the density itself. We therefore need to damp any errors we make in the density (also already for much smaller density changes), or a strong artificial potential is often generated. We note that this feature is much less severe for dipole tracking [16], where almost any naive implementation will perform adequately. However, in the case of the density inversion much more caution is needed. This also applies to the calculation of the initial state which needs to be accomplished with high accuracy even

[1] This remains true also for more advanced Runge–Kutta methods and even the Krylov subspace based Lanczos method [29] and therefore a more fundamental approach to cure this problem is needed.

in the low density regions. The way we ensure stability and accuracy of the inversion method is explained in detail in the next sections, while readers only interested in the basic ideas may skip to the examples Section 7 right after this section.

Note that tracking can also be used for other mutually conjugate variables like the potential and density. This can be used to construct exact functionals also for other types of TDDFT like time-dependent current or dipole DFT, or TDDFT for cavity quantum electrodynamics [30]. For example, consider the potential $v(\mathbf{r}t) = v_0(\mathbf{r}t) + \varepsilon(t)x$, where $v_0(\mathbf{r}t)$ is fixed and usually time-independent, e.g., it could be the potential from the nuclei in some molecule. In contrast $\varepsilon(t)$ is the field strength of a linearly polarised laser (in the x direction treated in dipole approximation), that one can adjust. In this case the field strength $\varepsilon(t)$ and dipole moment $\mu_x(t) = \langle\Psi(t)|\hat{\mu}_x|\Psi(t)\rangle$, $\hat{\mu}_x = \sum_{i=1}^{N} x_i$, are also conjugates [31], and one can use tracking to find the field strength $\varepsilon(t)$ that from a given initial state $|\Psi_0(t)\rangle$ yields a specified dipole moment $\mu_x(t)$ exactly the same way as we find the potential $v(\mathbf{r}t)$ that yields a specified density $n(\mathbf{r}t)$. Instead of equation (8) we just take the second time-derivative of $\mu_x(t)$, instead of $n(\mathbf{r}t)$, to find a similar relation for this case [16].[2]

3 Numerical framework

In this section we describe a time-propagation method that will allow us to solve the TDSE much more efficiently (as it allows for very long time-steps) and accurately than with the Euler method. To describe the method we again discretize the time-axis into small pieces of length Δt. The time-stepping for a time-dependent Hamiltonian $\hat{H}(t)$ is presented pictorially as follows[3]

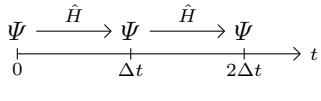

in which every arrow denotes a time evolution of the wave function on each interval using a Hamiltonian $\hat{H}$ obtained from $\hat{H}(t)$ by evaluating it at the midpoint of each time-interval and employing it for the whole of each of the intervals Δt. Since the only time-dependent part of the Hamiltonian is the potential this requires the evaluation of the potential at the midpoint of each interval. The time-propagation on each interval is performed with the evolution operator $\hat{U}(\Delta t) = e^{-\mathrm{i}\hat{H}\Delta t}$ which we will henceforth denote as the time-step operator. It can be shown rigorously [28,32,33] that for Lipschitz continuous potentials we will converge to the exact continuous solution of the TDSE in the limit that Δt approaches zero. This procedure allows us to step through time with a computational cost which is determined by the fact that we have to apply $\hat{U}(\Delta t)$ repeatedly. As there exist good efficient approximations to this operator, this approach is both efficient and accurate.[4] Note that the use of midpoint potentials has the important property that it ensures time-reversal symmetry and allows for longer steps than for other choices of time-points to evaluate the potential in.

The main question that remains is how to apply the formal time-step operator $\hat{U}(\Delta t) = e^{-\mathrm{i}\hat{H}\Delta t}$. A mathematically well-defined way of doing this is to first compute all the eigenstates $|n\rangle$ and energies E_n of the Hamiltonian $\hat{H}$ and then expand the state $|\Psi\rangle$ in terms of these eigenstates, and use that the time-evolution of each eigenstate is given by a simple exponential, i.e.,

$$\hat{U}(\Delta t) = \sum_n e^{-\mathrm{i}E_n\Delta t}|n\rangle\langle n|. \tag{9}$$

It is, however, generally very expensive to compute all the eigenstates and this approach is therefore not very practical, especially since the operator $\hat{H}$ changes every time step. Instead, one usually assumes that the formal time-step operator may be Taylor expanded,

$$\hat{U}(\Delta t) = \sum_n \tfrac{(-\mathrm{i})^n}{n!}(\Delta t\hat{H})^n, \tag{10}$$

since this sum can be truncated at low order while keeping a high precision thanks to the rapid decay of the terms $(\Delta t)^n$ for small Δt. We will use the Krylov subspace projection method outlined in [29], which is also known as the Lanczos method [24], to evaluate $\hat{U}(\Delta T)|\Psi\rangle$ approximately. This method can be viewed as a way to get parts of all higher order terms in the truncated sum in equation (10) at no extra significant cost.[5] The approach used here is more advantageous than general purpose solvers like the Runge–Kutta methods, since it has been made specifically for the TDSE. The exact time-step operator preserves the norm exactly, as well as the time-reversal symmetry since we use midpoint Hamiltonians. The Krylov subspace projection method [29] that we use closely approximates the exact operator, and therefore also keeps these properties to a high extent, while the Runge–Kutta method does not build them in. We further note that in equation (10) we usually compute $\hat{H}|\Psi(t)\rangle$ by direct application of the operator instead of using the

[2] The conjugate pairs are very special though. For example, say we start in a ground state, and prescribe that $S(t) = \langle\Psi(t)|\hat{O}|\Psi(t)\rangle$ should go from 0 to 1, with $S(t)$ the occupation of the first excited state, $\hat{O} = |\psi_1\rangle\langle\psi_1|$. That is, we want to find the field strength $\varepsilon(t)$ that takes us from the ground state to first excited state in some specific manner. Then most likely there just exists no such $\varepsilon(t)$, as most tracks are simply not physically possible, and so tracking will usually run into infinite field strengths $\varepsilon(t)$ [16,17]. In contrast, for $\mu_x(t)$ TDDFT guarantees us the existence of a unique $\varepsilon(t)$, so in this case, or for other conjugates, we never run into trouble.

[3] For simplicity we employ an equidistant time grid but all time-stepping strategies that we present, including those for inversions, can also use adaptive time steps to boost performance.

[4] For Hamiltonians that change rapidly in time it can be advantageous to allow for temporal change of the Hamiltonian within each interval by using the so called Magnus expansions [29] but usually this is not needed.

[5] Presently the Krylov subspace projection method is generally the method of choice although for wave functions with a very broad spectrum the split-operator method [24] may perform better.

eigenvector expansion, which is considerably more efficient than using equation (9). For example, in real space, we can directly evaluate the appropriate derivatives for $\hat{T}$ and do the multiplications with $\hat{W}$ and $\hat{V}(t)$ to apply $\hat{H}(t)$ to a wave function.

We finally like to point out a fundamental difference between equations (9) and (10), which will be relevant for the discussion of zero boundary conditions. The difference only exists in the spatial continuum limit in which we have to deal with unbounded operators (for a more detailed mathematical discussion we refer to [20]). Our discussion of the time-stepping above was completely independent of the spatial representation and in particular the operator $\hat{U}(\Delta T)$ as defined by equation (9) is well-defined in the continuum limit on the domain of wave functions on which the Hamiltonian itself is defined. However, the expansion in equation (10) is defined on a much smaller domain [23] and only allows for (spatially) analytic potentials and initial states since only then one can apply the Hamiltonian an arbitrary number of times. In contrast, equation (9) applies also to systems that contain potentials with spatial jumps, such as the finite square well. Another important difference involves the incorporation of boundary conditions. While for equation (9) the spatial boundary conditions on the system are contained in the eigenstates $|\psi_n\rangle$, the right hand side of equation (10) at first sight seems not to contain any information on the boundary conditions of the system. The only way to incorporate them is in the specification of the domain of the operators $\hat{H}^n$ for all n [23]. Since we prefer to use equation (10) for practical purposes it is important to know how this is done. If an initial wave function, for example, has periodic or hard-wall boundary conditions we like the operator $\hat{H}^n$ to preserve that property. How this is dealt with in practice is explained in detail in Appendix A.

The discussion up to now only involved the temporal aspect of our framework. The remaining issue is to find a suitable spatial representation to use. For general time-dependent systems the wave function and density can vary considerably both in space and time, and are difficult to represent in terms of (time-independent) basis functions (in contrast to for weakly perturbed systems). It is therefore expedient to use an equidistant grid, as it will essentially (dependent on the boundary conditions) represent the wave function equally well independent of its spatial localization. For this reason real space codes are very common in the study of strong field dynamics and we will adopt the same strategy. An alternative would be to use time-dependent basis functions. To evaluate the necessary spatial derivatives we will further make use of common finite difference techniques but make them more accurate at the boundaries.

4 Low density regions

In this section we demonstrate why it is important to compute the density accurately also in the low density regions and how to do this. To illustrate the problem it suffices to consider a single spin-free electron in 1D. We consider the Hamiltonian $\hat{H} = -\frac{1}{2}\partial_x^2 - \cos(\frac{2\pi x}{L})$ with periodic

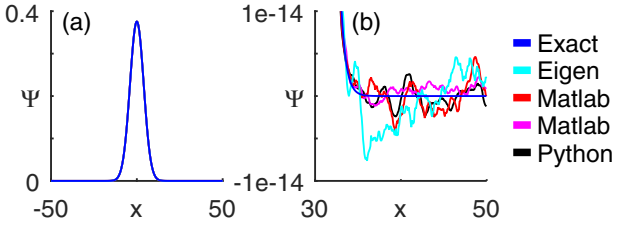

Fig. 1. The ground state of the potential $v(x) = -\cos(\frac{2\pi x}{L})$ with $L = 100$, (a) in the whole interval and (b) zoomed in on the vanishing density region, for five different eigensolvers. An exact eigensolver, designed to also yield the correct result in vanishing density regions, Eigen, a C++ library eigensolver, Matlab's eig and eigs solvers, and Python's eigsh. Only the exact eigensolver is correct for vanishing densities.

boundary conditions and take $L = 100$. By employing a variety of different eigensolvers we can calculate the ground state. A sample of such calculations is presented in Figure 1a. If we zoom in on the vanishing density region, see Figure 1b, we see that all the library eigensolvers produce a very irregular wave function in the region of function values of around 10^{-14}. Since we know the exact wave function (which is a Mathieu function) to be smooth and nodeless these irregularities merely represent numerical noise. The errors in this region are very large as the values are off by as much as 13 orders of magnitude and a randomly oscillating sign. Library eigensolvers by design usually only guarantee that the eigenvalues are converged to machine precision, while they make no further guarantees for the precision of the eigenstates. For many purposes, such as in the calculation of expectation values, such errors are generally negligible in the quantities that are being calculated. However, for the inversion problems that we are concerned with in this work such noise deserves careful attention even if the densities are low. This is again illustrated with our example of the ground state of Figure 1 of which we take the square to obtain the density. Suppose we now want to find the potential that yields this density by means of a numerical inversion method. This should give back the $-\cos(\frac{2\pi x}{L})$ potential, and indeed it does provided we use the exact ground state density. However, if we use one of the noisy ground states to construct the density, we can only expect to obtain an accurate potential in the inner region where the ground state is well represented, while in the outer region we will be trying to find a potential that yields very irregular noise. Therefore the algorithm to find the potential will likely not converge in this region, although this can easily be fixed by for example imposing a suitable artificial form of the potential in the outer region.

In the time-dependent case the situation is much more intricate as any noise in the outer region tends to propagate quickly into the higher density regions as we step through time. One can still handle many simple cases by artificially modifying the potential in the outer region [34]. However, this will cause the density we actually get to drift apart from the prescribed density. This is an issue, since tracking by design tries to get to the prescribed density within each time-step. Our method will therefore try to compensate the full difference within each time-step,

which requires a stronger and stronger potential, and one is again limited by the time it takes the errors to propagate inwards and grow in size, although much slower. This raises the question how to construct a precise black box approach to avoid excessive numerical noise and its propagation in time-dependent inversions. The first step is to compute the initial state to a high accuracy also in the low density regions. If the initial state is given, for example in some analytic form, then this is a straightforward procedure. However, usually the initial state is required to be an eigenstate of some Hamiltonian and we therefore need to solve the time-independent Schrödinger equation which is usually solved iteratively, like other large eigenvalue problems. The way to achieve high accuracy is to use an extremely strict convergence criterion and a method able to reach such high accuracy as discussed further in Appendix B. The other issue is the propagation of numerical errors. However, if the initial state is determined very accurately both propagation and time-dependent inversions can be done to high accuracy as well with little extra computational cost.

Now, in most situations we do not have to deal with such vanishing densities if we use a small domain. However, we stress that it is equally important to also treat the moderately low density regions precisely, such as densities like 10^{-4}, which we commonly encounter. Otherwise the same issues as described in this section arise. The most important message of this section is thus that if we use any approximations, whether it is in prescribing the initial state $|\Psi_0\rangle$ and density $n(\mathbf{r}t)$ or anywhere in the inversion method, these approximations have to also be reasonable in the moderately low density regions. The approximations do not necessarily have to be precise, but they should at least be reasonable and consistent enough not to cause too wildly fluctuating potentials during the inversion as this will cause a breakdown of the method. This prohibits the use of many common approximations, which have often only been designed to get the high density regions right while even moderately small density regions may again be very non-physical. One is forced to pay careful attention to the physics in the moderately low density regions. Initial states also always have to be precise (possible within a certain approximation). In particular many quantum codes designed for large systems by default only treat the high density regions accurately (to save computational effort), and therefore one should be careful if using initial states from such codes.

5 Time-independent inversion

After the preliminaries in the previous sections we proceed with the discussion of the actual inversion methods. In this section we consider the case of a time-independent system and look for a time-independent potential $v(\mathbf{r})$ which produces a given ground state density $n(\mathbf{r})$.[6] A simple way to find this potential is to use an iterative method along the lines of [25]. We define an iterative sequence $v_k(\mathbf{r})$ of potentials by the solution of

$$v_{k+1}(\mathbf{r}) - v_k(\mathbf{r}) = \delta \ln\left(\frac{n_k(\mathbf{r})}{n(\mathbf{r})}\right), \tag{11}$$

where $n_k(\mathbf{r})$ is the ground state density obtained by solving the Schrödinger equation with potential $v_k(\mathbf{r})$, δ a positive number, and $n(\mathbf{r})$ the desired density. We repeatedly compute the ground state density until it converges, $n_k(\mathbf{r}) \to n(\mathbf{r})$, while updating $v_k(\mathbf{r})$ each time. Note that per construction, the procedure is guaranteed to yield the right result if it converges, as the source term (the right hand side of Eq. (11)) only vanishes if $n_k(\mathbf{r}) \to n(\mathbf{r})$. We further see that the procedure in each iteration increases the potential where the density $n_k(\mathbf{r})$ is too large, and thereby reduces the density in that region, and decreases the potential where the density $n_k(\mathbf{r})$ is too small, to attract more density. This method performs equally well for interacting and non-interacting systems, except that it is usually much faster in the latter case, as one can usually employ orbitals in that situation.

The main question remaining is how large a correction one should add in each iteration, to make it large enough that the procedure converges fast, yet small enough not to over-correct so strongly that it causes the potential to fluctuate wildly. The logarithmic right hand side of equation (11) represents a decent first guess, since even if one starts with a completely wrong potential $v_0(\mathbf{r})$ for which the density $n_0(\mathbf{r})$ is off by orders of magnitudes at some points, it still causes only a limited correction in the right direction (even for a quite large δ). When the density is nearly correct it causes a large enough change in the potential to ensure a reasonably fast convergence. This approach is never very fast though because the local corrections need to gradually shape the global potential that we seek for, but many other methods share this feature. There are many options for improvement, like an adaptive $\mathbf{r}$-dependent δ, but not an easy answer to what is best in general and a more in-depth discussion is beyond the scope of the present work.

If one imposes boundary conditions on the wave function that are consistent with the prescribed density, and is cautious with the low density regions as described in the previous section (or enforce an artificial potential in the low density regions), the above method usually works trivially. There also exists a plethora of other methods to compute $v[n](\mathbf{r})$ (see, e.g., [26]), some of which may be somewhat faster than the above method, but many of these methods often require extra care to perform, and/or have other limitations. Many of them are further specialised to non-interacting systems and orbitals for a slight performance gain at the cost of generality, as such a specialisation can be optimized further than a generic algorithm. So the above method is a safe generic choice.

6 Time-dependent inversion

First we develop a more precise time-stepping strategy to do the tracking. The key ingredient in this strategy is

[6] This also lets us compute $v_{\mathrm{Hxc}}[n] = v_s[n] - v[n]$ by applying the method to both a non-interacting and interacting system and subtracting the two results.

an update formula that we obtain by recasting a recent fixed-point approach to TDDFT [18,19] in the form of this more precise time-stepping strategy. These two items combined form our final inversion algorithm.

6.1 Time-stepping strategies

We recall that in the time-stepping strategy of Section 3 we needed the potential at the midpoints of the time-intervals and will in this section describe several ways to compute this given the density $n(\mathbf{r}t)$ instead of the potential $v(\mathbf{r}t)$.

We will in the following denote the potential at the nth midpoint by $\bar{v}(\mathbf{r}t_n) = v(\mathbf{r}, (t_n + t_{n-1})/2)$. Our goal is to evaluate this quantity and we assume that we have obtained the potential $\bar{v}(\mathbf{r}t_m)$ for $m < n$ at the midpoints of the previous time intervals. We further aim to set up an iterative procedure to obtain successively better potentials $\bar{v}_k(\mathbf{r}t_n)$ which will converge to $\bar{v}(\mathbf{r}t_n)$ if we increase k. We may get a good estimate for the initial guess $\bar{v}_0(\mathbf{r}t_n)$ at low computational cost by extrapolation. For example, we may use linear extrapolation $\bar{v}_0(\mathbf{r}t_n) = 2\bar{v}(\mathbf{r}t_{n-1}) - \bar{v}(\mathbf{r}t_{n-2})$, or even simpler $\bar{v}_0(\mathbf{r}t_n) = \bar{v}(\mathbf{r}t_{n-1})$, although in practice we recommend the use of higher orders as the computational cost is very low. It is also possible to compute the previous on-point potentials $v(\mathbf{r}t_m)$ for $m < n$ by solution of equation (8) at time t_m (since we know the wave functions $|\Psi(t_m)\rangle$, see Sect. 2.2) and use these in the extrapolation as well. This is most useful in the first step where we only have access to $v(\mathbf{r}t_0)$ to do our extrapolation. Once we have obtained the extrapolation $\bar{v}_0(\mathbf{r}t_n)$ we often want to improve it by some corrector steps before we use it to take a next time step.

One common way to correct the midpoint potentials $\bar{v}_k(\mathbf{r}t_n)$ is to use it in a time propagation step to obtain $|\Psi_k(t_n)\rangle$ which approximates $|\Psi(t_n)\rangle$. We subsequently use $|\Psi_k(t_n)\rangle$ to obtain an approximate on-point potential $v_k(\mathbf{r}t_n)$ and construct a better midpoint potential using $\bar{v}_{k+1}(\mathbf{r}t_n) = \frac{1}{2}(v(\mathbf{r}t_{n-1}) + v_k(\mathbf{r}t_n))$.[7] This approach is often used in KS theory where the use of two corrector steps usually works well [2].[8] Another common choice in KS theory is to use no corrections at all and to employ a higher order extrapolation instead at the expense of requiring shorter time steps. This approach performs well for dipole tracking but is generally not accurate enough in our case.

To achieve much greater stability we use for our inversion method a different way to update the potential which uses the control target that we want to obtain. Again we use $\bar{v}_k(\mathbf{r}t_n)$ to perform a time-step to get $|\Psi_k(t_n)\rangle$ and a density $n_k(\mathbf{r}t_n)$. Subsequently we compute the residual $n_k(\mathbf{r}t_n) - n(\mathbf{r}t_n)$ to measure how different the density that we obtained is from the density $n(\mathbf{r}t_n)$ that we want to produce. The idea is then to find an update formula of the form

$$\bar{v}_{k+1}(\mathbf{r}t_n) = \bar{v}_k(\mathbf{r}t_n) + f[n_k(\mathbf{r}t_n) - n(\mathbf{r}t_n)], \tag{12}$$

where f is a given functional of the residual $n_k - n$ at time t_n that corrects the midpoint potential in such a way that $\bar{v}_{k+1}(\mathbf{r}t_n)$ yields a significantly smaller residual. The detailed form of this functional is presented in the next section as it is irrelevant for the discussion of the general update strategy here. We use it to repeat the update in equation (12) a few times until the residual becomes very small before we move to the next time-step. The large advantage of this strategy is that it always tries to compensate any tiny error in the time-propagation by marginally changing the midpoint potentials to ensure that the target density still is produced very accurately. Even if $n(\mathbf{r}t)$ changes by orders of magnitude this procedure will still produce the density correctly (while the potential will also be very accurate as we stay on the prescribed density path in time).

In contrast, the standard correction formulas for the potential that we mentioned before do not correct any error in $n(\mathbf{r}t)$ except by sheer luck, and for example what may be a small error to start with may become a major error if $n(\mathbf{r}t)$ decreases over time by orders of magnitude, which causes a breakdown of the method. In short, our approach may be slightly more expensive but is much more stable. We shall also see in Section 6.2 that it also allows for further stabilising techniques.

6.2 Update formula

We now address the question how to derive a suitable update procedure that can be used in equation (12). To do so we use equation (8) to define an iterative sequence v_k of potentials on a given time interval, starting from an initial guess v_0, by the solution of

$$-\nabla \cdot (n(\mathbf{r}t)\nabla v_{k+1}(\mathbf{r}t)) = q[v_k](\mathbf{r}t) - \partial_t^2 n(\mathbf{r}t). \tag{13}$$

In previous works [18,19] we showed (under mild assumptions) that indeed $v_k \to v[\Psi_0, n]$ in Banach norm sense. We will reformulate this procedure in a way that suits our time-stepping strategy. For much greater efficiency we therefore compute the potential on many successive small time-intervals of length Δt instead of on a long global time interval.[9] We further find it advantageous to eliminate the quantity $q[v_k]$ from equation (13) to obtain a simpler equation that depends on densities and potentials only. This can be done by employing equation (8) for a system with potential v_k. We then obtain

$$\begin{aligned}-\nabla \cdot (n(\mathbf{r}t)\nabla v_{k+1}(\mathbf{r}t)) = \partial_t^2 \left[n[v_k](\mathbf{r}t) - n(\mathbf{r}t)\right] \\ - \nabla \cdot (n[v_k](\mathbf{r}t)\nabla v_k(\mathbf{r}t))\,.\end{aligned}$$

Furthermore, since $n[v_k]$ approaches n as v_k approaches $v[\Psi_0, n]$ we can replace the last $n[v_k]$ by n close to

[7] It is important to add the potentials but not the wave functions to obtain $\bar{v}_{k+1}(\mathbf{r}t_n)$ since adding wave functions does not conserve the norm.

[8] This approach applies equally well to KS theory to compute $v_{\mathrm{Hxc}}(\mathbf{r}t_n)$ and thus the midpoint Hamiltonian from $|\Phi(t_n)\rangle$ (and the already known $|\Phi(t_m)\rangle$ if we include memory).

[9] One can only expect to find $v[\Psi_0, n]$ accurately at a given time if it is already correct at previous times.

convergence (which we always are as we find $v[\Psi_0, n]$ step-by-step).[10] We then obtain

$$-\nabla\cdot(n(\mathbf{r}t)\nabla\left[v_{k+1}(\mathbf{r}t)-v_k(\mathbf{r}t)\right])=\partial_t^2\left[n[v_k](\mathbf{r}t)-n(\mathbf{r}t)\right]. \quad (14)$$

This equation shows us how we can update v_k given the residual $n[v_k]-n$, and (once time-discretized) provides our desired update formula. The functional f in equation (12) is therefore implicitly defined by solution of (14). We can however improve the update formula significantly by also making use of the current. If we use the continuity equation (6) we may also write

$$\nabla\cdot(n(\mathbf{r}t)\nabla[v_{k+1}(\mathbf{r}t)-v_k(\mathbf{r}t)])=\partial_t[\nabla\cdot\mathbf{j}[v_k](\mathbf{r}t)+\partial_t n(\mathbf{r}t)]. \quad (15)$$

While equation (14) is only converged if v_k yields n, equation (15) is instead converged when v_k yields a current density $\mathbf{j}[v_k]$, that together with n satisfy the continuity equation (since exactly then does the source terms on the right hand side vanish). By combining these two equations we can require both n and $\nabla\cdot\mathbf{j}$ to be as accurate as possible, even when numerical errors try to build up. This turns out to be another crucial ingredient in stabilising the numerics and is perhaps the most crucial innovation as it also acts as a damper on all errors.[11] We use a parameter μ (chosen as we see fit) to weigh the two equations. This results in

$$\begin{aligned}-\nabla\cdot&\left(n(\mathbf{r}t)\nabla\left[v_{k+1}(\mathbf{r}t)-v_k(\mathbf{r}t)\right]\right)\\&=(1-\mu)\partial_t^2\left[n[v_k](\mathbf{r}t)-n(\mathbf{r}t)\right]\\&\quad-\mu\partial_t\left[\nabla\cdot\mathbf{j}[v_k](\mathbf{r}t)+\partial_t n(\mathbf{r}t)\right].\end{aligned} \quad (16)$$

The only thing left to do now in order to arrive at the desired update formula is to discretize this formula with respect to time. By using only times t_m with $m\leq n$ for the derivatives, and employing the fact that we have already converged the density and current up to time t_{n-1}, we obtain

$$\begin{aligned}-\nabla\cdot&\left(\bar{n}(\mathbf{r}t_n)\nabla\left[\bar{v}_{k+1}(\mathbf{r}t_n)-\bar{v}_k(\mathbf{r}t_n)\right]\right)\Delta t^2\\&=A\left[n[v_k](\mathbf{r}t_n)-n(\mathbf{r}t_n)\right]\\&\quad-B\Delta t\left[\nabla\cdot\mathbf{j}[v_k](\mathbf{r}t_n)+\partial_t n(\mathbf{r}t_n)\right],\end{aligned} \quad (17)$$

as our update formula. The constants A and B depend on the discretization scheme and on the parameter μ of equation (16), which effectively leaves the choice of their values at our disposal. In practice we usually employ values between 0.5 and 1 for them.[12]

Note that there exist very efficient numerical methods to solve the Sturm–Liouville problem of equation (17), i.e., to invert the self-adjoint operator $-\nabla\cdot(n(\mathbf{r})\nabla)$ efficiently. The solution is determined up to a constant which is fixed by a gauge choice (for instance, we can require the spatial integral of the potential to be zero). In the real space finite difference framework that we used in our work relaxation methods are a rather slow, but easy to implement, choice for this inversion [35]. For greater speed one can accelerate these methods using the significantly more involved multi-grid methods also given in [35]. This generally leaves the computational cost negligible compared with the one required in the time-propagation of the wave function.[13] As iterating equation (17) further generally converges within 5 to 10 iterations for each time-step, our algorithm is essentially as fast as the time-stepping scheme itself. We also find that the precision of the resulting potential is limited mainly by the time-step operator (assuming a sufficient spatial resolution). By increasing the precision thereof almost arbitrary precision can be achieved even when the density changes by many orders of magnitude.

Finally, note that our new update formula equation (17) is a considerable improvement over that of our prior work (see Eq. (10) in [10]). In that work we only considered 1D multi-particle problems where a direct integration of the Sturm–Liouville operator is possible. However, solving equation (17) directly, instead of by using two anti-derivatives, allows us to treat 2D and 3D systems. Moreover, it is also simpler since we bypass a non-trivial integration constant.[14] Finding a very precise anti-derivative is often as difficult as inverting the Sturm–Liouville operator so although it may seem easier at first glance we discourage this route. The new update

[10] The true right hand side of equation (14) has an extra term $-\nabla\cdot((n[v_k]-n)\nabla v_k)$ that vanishes at convergence. Removing it therefore has no effect on the final result if we converge. It also does not affect the convergence rate itself as it only adds $-\nabla\cdot((\bar{n}[v_k](\mathbf{r}t_n)-\bar{n}(\mathbf{r}t_n))\nabla\bar{v}_k(\mathbf{r}t_n))\Delta t^2$ to the right hand side of equation (17). If the potentials only differ in the last time step the size of this term is of the order Δt^4 while the size of all the terms in equation (17) are of the order Δt^2.

[11] In simple terms it ensures that both the modulus and phase of the wave function remain accurate. Adding more conditions for even greater stability is possible but these conditions will have to be of a consistent form and have to perform well in combination. This is especially relevant to extent the method to other conjugate pairs.

[12] The scale of A and B determines the scale of the correction $\bar{v}_{k+1}(\mathbf{r}t_n)-\bar{v}_k(\mathbf{r}t_n)$ that we add per iteration and should be large enough to converge fast yet small enough to not over-correct. Different relative weights of A and B eliminate different error build ups slightly better, and it is mainly important to have some of each. The presented values generally perform well and when the choice matters it is usually due to an insufficient grid or time-propagation method.

[13] Especially with multi-grid acceleration one should be cautious to use a proper discretization of the Sturm–Liouville operator. For example in 2D a proper 5-point formula is $\frac{1}{2\Delta x^2}[(n_{i-1,j}+n_{i,j})v_{i-1,j}-(n_{i-1,j}+2n_{i,j}+n_{i+1,j})v_{i,j}+(n_{i,j}+n_{i+1,j})v_{i+1,j}]+\frac{1}{2\Delta y^2}[(n_{i,j-1}+n_{i,j})v_{i,j-1}-(n_{i,j-1}+2n_{i,j}+n_{i,j+1})v_{i,j}+(n_{i,j}+n_{i,j+1})v_{i,j+1}]$. This preserves the self-adjointness of the continuous Sturm–Liouville operator. Crucial to the stability of the multi-grid method [35], once divided by its diagonal, the eigenvalues further lie from 0 to 2 with only one zero eigenvalue (for v constant). For increased efficiency one may sometimes further weigh the two directions when dividing by the diagonal. Within the multi-grid approach the computational cost of inverting a Sturm–Liouville operator is only a few times that of applying the Sturm–Liouville operator itself, so it is small compared to the cost of taking a time-step for the time-propagation of the wave function (except in the one-electron case or for a few orbitals).

[14] The first integration of equation (17) gives rise to a boundary condition dependent integration constant, while the integration constant of the second integration (and of solving Eq. (17) directly) is the gauge constant of the potential. The first integration constant is for the even-odd boundary conditions (see Appendix A) fixed by the requirement that the anti-derivative must be odd. On $\mathbb{R}$, the anti-derivative must also vanish at $\pm\infty$ (otherwise the potential we obtain will approach $\pm\infty$ for $\pm\infty$ extremely fast). For periodic boundary conditions the integration constant is fixed by periodicity [10].

formula also has the advantage that we no longer need to prescribe $\mathbf{j}(\mathbf{r}t)$ or to smoothen the potential.[15]

We finally like to remark that there exists a special time-dependent one-orbital inversion formula [6,36] which is very valuable for comparisons and testing of our numerical inversion method, especially since there are multiple ingredients necessary for our method to work properly. Since the physics of the low density regions can be intricate it may be worth to first recreate some of our examples or to compare to this one-electron formula for some very simple cases when designing a computer code.

7 Examples

To illustrate the capabilities of our procedure we apply it to small interacting systems as well as larger non-interacting systems and use it to construct the exact Hxc potentials for small systems. We first apply it to a simple 1D model system in Section 7.1 to illustrate the main aspects of the method after which we consider 2D and 3D model systems in Sections 7.2 and 7.3 to confirm that the method works the same in 2D and 3D.

We note that one can also easily do more realistic cases like using (soft) Coulomb potentials. Indeed, we have already performed many such computations. We have also used zero boundary conditions, handled density changes of up to 34 orders of magnitude, used potentials and interactions with kinks (continuous but non-differentiable functions) and much more. We find it educational to concentrate on one simple model here in order to focus on the method rather than on the systems and to provide other examples in a follow up paper. For some examples with soft Coulomb and zero boundary conditions see also [37], for examples involving scattering processes see [38,39], and for an example in the context of cavity quantum electrodynamics see [30].

Also note that all computations were done on a laptop for which the cases of two interacting electrons in 2D or 500 non-interacting electrons in 3D are currently about the limit for a simple real space code. When using distributed memory two interacting electrons in 3D also becomes possible. Furthermore, one can handle much larger systems by using approximate representations which, if precise enough, could still yield nearly exact results.

The computed potentials of our examples were checked to closely reproduce the prescribed density and its time-derivative (by design of the method). We also verified that they are independent of the size of the time step and the operator as well as the spatial representation. Some of the potentials were also tested by means of the analytic formula available for the single orbital case.

[15] In 1D we can determine the current $\mathbf{j}(\mathbf{r}t)$ up to a constant for a given density $n(\mathbf{r}t)$ using the continuity equation (6). The 3D analogue of this procedure amounts to the determination of the longitudinal current which can be done by using relaxation methods. However, by using $\partial_t n(\mathbf{r}t_n)$ via the continuity equation we do not even need to prescribe $\mathbf{j}(\mathbf{r}t)$ at all.

7.1 1D model system

In this part we consider N electrons on a quantum ring of length $L = 10$, over a time period of length $T = 20$. We start by computing an initial state, $|\Psi_0\rangle$, which in all cases is a ground or excited state of a (properly periodic) Hamiltonian with external potential v_0 and interaction w given by

$$v_0(x) = -\cos\left(\tfrac{2\pi x}{L}\right),$$
$$w(x_1, x_2) = \lambda \cos\left(\tfrac{2\pi(x_1 - x_2)}{L}\right),$$

where λ is the interaction strength. We then construct the corresponding initial density $n_0(x)$, and from it the (spatially periodic) time-dependent densities n_1 and n_2:

$$n_1(xt) = n_0(x - r(t)), \tag{18}$$
$$n_2(xt) = \tfrac{1}{2}[n_0(x - r(t)) + n_0(x + r(t))], \tag{19}$$
$$r(t) = \tfrac{L}{2}\left[1 - \cos\left(\tfrac{\pi t}{T}\right)\right]. \tag{20}$$

The density n_1 describes a situation where the initial density n_0 is rigidly translated around the ring exactly once, while n_2 describes a situation where the initial density n_0 is split in equal halves that are rigidly translated in opposite directions to rejoin at times $\frac{T}{2}$ and T.[16] We used our algorithm to obtain the potentials that produce these prescribed densities n_1 and n_2 via time-propagation of the initial state $|\Psi_0\rangle$ by the TDSE. This was done for interaction strengths $\lambda = 1$ and $\lambda = 0$, which we will in the following respectively call the interacting and non-interacting case. It is worth noting that $v(\mathbf{r}t_0) = v_0(\mathbf{r})$ only when $\partial_t^2 n(\mathbf{r}t_0) = 0$ if $|\Psi_0\rangle$ is an eigenstate.[17] If we want $v(\mathbf{r}t_0) = v_0(\mathbf{r})$ we need to prescribe a density with $\partial_t^2 n(\mathbf{r}t_0) = 0$. This is the case for the splitting densities but not for the translated ones. Our examples are made to be numerically challenging since the external potential has to change the density by many orders of magnitude (for a fixed spatial coordinate), while preserving accuracy. In the case of the density splitting the external potential further has to ensure that the wave function splits and recombines correctly.

7.1.1 One-electron

We consider the case of a single electron and take the initial state to be one of the (spatially identical) degenerate spin-up or spin-down ground states. The corresponding potentials and densities (insets) are shown in Figure 2. As we have just argued above, $v(\mathbf{r}t_0) = v_0(\mathbf{r})$ for the split density, but not for the translated one. Moreover the potentials have some additional properties in this simple case as we see they are periodic in time, with period $2T$ for the translated and T for the split density. The periodicity

[16] In practice we use cubic interpolation to evaluate such densities, or bicubic and tricubic interpolations in 2D and 3D.

[17] As we start in a stationary state $|\Psi_0\rangle$, we have from equation (8) that $-\nabla \cdot (n(\mathbf{r}t_0)\nabla v_0(\mathbf{r})) = q(\mathbf{r}t_0)$. Subtracting this from equation (8) evaluated at t_0 yields $-\nabla\cdot(n(\mathbf{r}t_0)\nabla\,(v(\mathbf{r}t_0) - v_0(\mathbf{r}))) = -\partial_t^2 n(\mathbf{r}t_0)$, so indeed $v(\mathbf{r}t_0) = v_0(\mathbf{r})$ up to a gauge constant only if $\partial_t^2 n(\mathbf{r}t_0) = 0$.

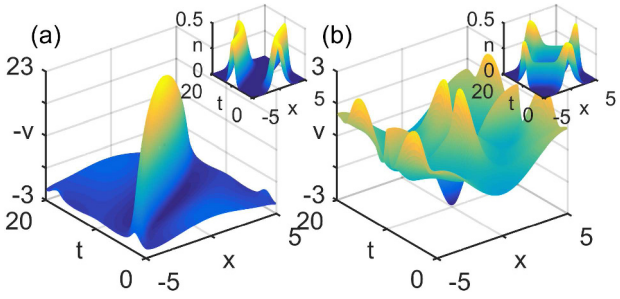

Fig. 2. The potentials that produce the (a) translated and (b) split one-electron densities (insets). Note that in (a) we plotted minus the potential for better visibility.

of the densities therefore carries over to the potentials (as do the other symmetries of the densities).[18] Note that for a translation we move in one direction around the ring from 0 to T and in the opposite direction from T to $2T$ to arrive at the initial position at time $2T$. For this simple case we even end up in the original state, i.e., the final state $|\Psi_T\rangle$ and initial state $|\Psi_0\rangle$ are the same. These properties are all special to the one-electron case, as we shall see next.[19] Finally note that especially in the translated case the potential changes most in the low density regions, which is a common feature of inversions.

7.1.2 Interacting electrons

Here we consider two interacting systems of two and three electrons respectively. In the first case, we take the initial state $|\Psi_0\rangle$ to be the two-electron singlet ground state (with a spatial part that is symmetric under exchange of the two spatial coordinates)[20] while in the second case we take the initial state to be the first stationary three-electron state with all spin up (or down) (with a spatial part which is anti-symmetric under interchange of the spatial coordinates). The corresponding potentials and densities (insets) are shown in Figure 3. As expected we again find that $v(\mathbf{r}t_0) = v_0(\mathbf{r})$ for the split densities, but not for the translated ones. In contrast to the one-electron case we now see that none of the potentials have any symmetry in time. Furthermore we never end up in the original state, i.e., $|\langle\Psi_T|\Psi_0\rangle|^2 < 1$ although the final density is again n_0 and the time-derivative of the density vanishes at this time. This illustrates that such symmetries are indeed special to the one-electron case (or perhaps also to more special multi-electron cases).

7.1.3 Non-interacting electrons

In this section we again consider the split densities and take the initial states to be the four non-interacting

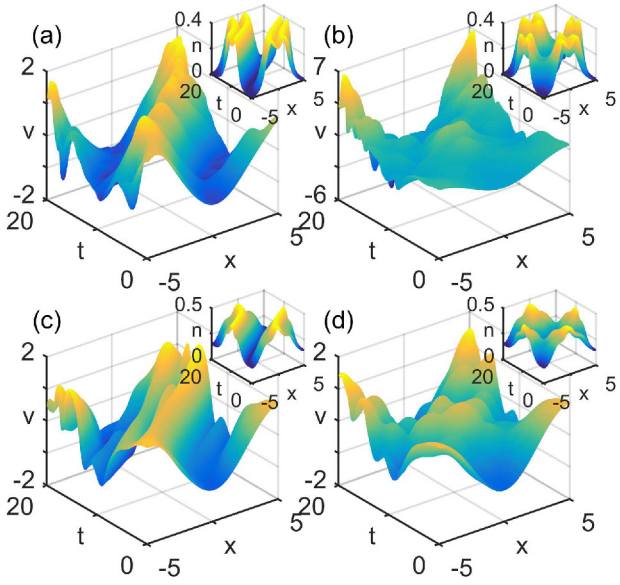

Fig. 3. The potentials that produce the (a) translated and (b) split interacting 2-electron densities (insets), and (c) translated and (d) split interacting 3-electron densities (insets).

2-, 6-, 10- and 14-electron ground states. We omit a discussion of the translated densities here, since we find it more interesting to show how much the potential changes character with the number of electrons for which the splitting case is sufficient. The many-particle initial states are constructed out of the orbitals presented in Figure 4. This allows us to treat large non-interacting systems since we can compute the ground states by computing the one-electron eigenstates and use them in the Aufbau principle, and we subsequently can employ our inversion procedure in terms of the one-electron orbitals.

In Figure 5, we show the potentials for a series of the four mentioned splits. The main feature that we see in Figure 5 is that the potentials become smoother the more particles we consider. We further note that $v(\mathbf{r}t_0) = v_0(\mathbf{r})$ for all the potentials, while only the two-electron potential is also periodic in time (as it is in fact equivalent with the single electron potential of Fig. 2b). Curiously, the 14 electron state is almost periodic again. However, for larger times (repeating the splittings) a more clear aperiodicity is expected to show up.

7.1.4 Benchmarking exchange-correlation approximations

Let us now give an example of one way to compute an exact Hxc potential and use it to benchmark the commonly used exact-exchange approximation [2]. For simplicity, we again consider the splitting of Figure 3b and consider our model system for two interacting electrons and the split density of equation (19).

To compute the exact Hxc potential we compute the interacting ground state $|\Psi_0\rangle$ and use it to construct the prescribed density $n_2(xt)$ and to find the exact KS initial state $|\Phi_0\rangle$ by a time-independent inversion. Then we use our method to compute the time-dependent potentials

[18] These symmetries are $v(x, \frac{T}{2}+t) = v(-x, \frac{T}{2}-t)$ in both cases while for the splitting further $v(x, \frac{T}{2}+t) = v(x, \frac{T}{2}-t)$.

[19] More precisely they apply to non-interacting one-orbital systems, of which a one-electron system is a special case. They also apply for two electrons in one orbital (see Sect. 7.1.3).

[20] We have already considered this simple example in [10]. However, since we use it later to benchmark Hxc approximations in Section 7.1.4 we present it also here.

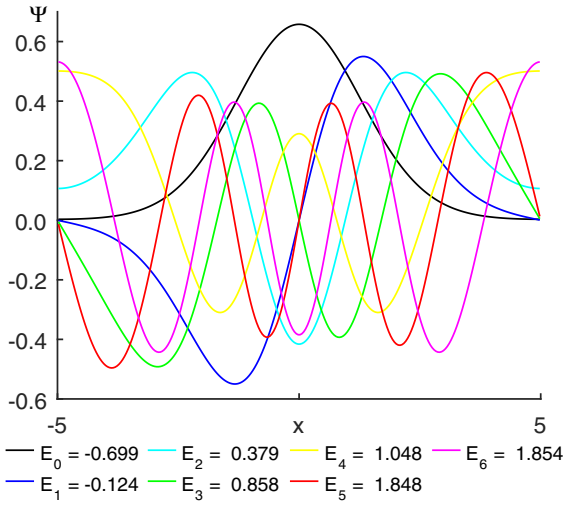

Fig. 4. The first seven one-electron eigenstates of $v_0(x)$.

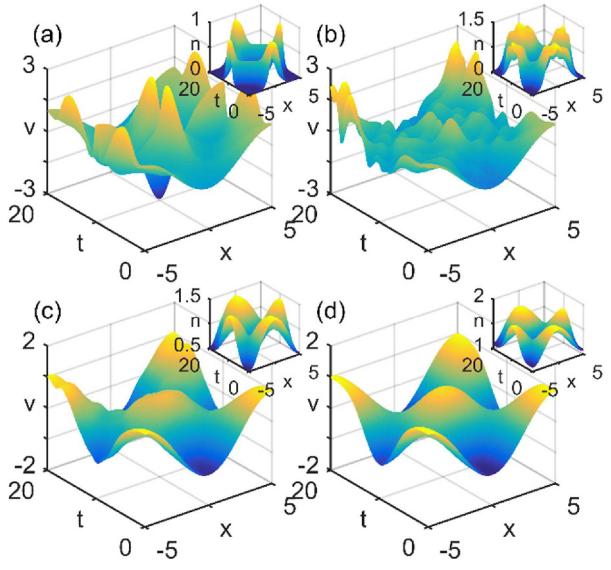

Fig. 5. The potentials that produce the split (a) 2-, (b) 6-, (c) 10- and (d) 14-electron densities (insets).

$v[\Psi_0, n_2]$ and $v_s[\Phi_0, n_2]$ and subtract these according to equation (5) to obtain v_{Hxc}. Note that $v[\Psi_0, n_2]$ is the potential that we already calculated to obtain Figure 3b.

Let us now consider the exact-exchange approximation for the Hxc potential. For two electrons this is particularly simple as it is equal to half the Hartree potential (the classical interaction contribution), $v_{\mathrm{Hx}}[n] = \frac{1}{2}v_{\mathrm{H}}[n]$ with $v_{\mathrm{H}}[n](\mathbf{r}t) = \int d\mathbf{r}' n(\mathbf{r}'t) w(|\mathbf{r} - \mathbf{r}'|)$. To evaluate this approximation we therefore only need to insert the exact density $n_2(xt)$ into the explicit expression for $v_{\mathrm{Hx}}[n]$. How this approximation compares with the exact Hxc potential is displayed in Figure 6.

We see from this figure that the approximation initially performs well but that it increasingly deviates more as time progresses. There is therefore a relation between the rapidity of the dynamics and the temporal length for

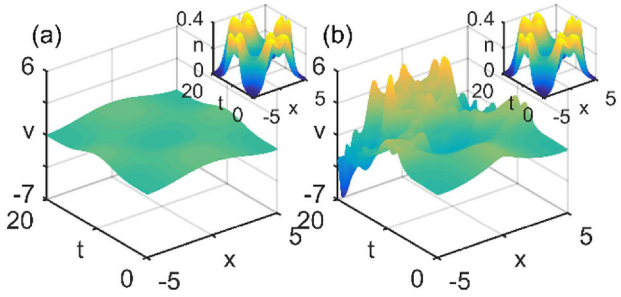

Fig. 6. (a) Approximate (exact-exchange) and (b) exact Hxc potentials for the exact split 2-electron density (insets).

which the approximation stays accurate. As memory is an important ingredient of the exact Hxc potential an adiabatic approximation depending on the instantaneous density (of which exact-exchange is a special case) will not reproduce the exact potential at later times. For our example an adiabatic approximation will always be periodic and symmetric around $t = 10$ in contrast with the exact potential. Therefore, if we want to use approximate functionals to study more rapid density variations, or dynamics over longer times, then we need to go beyond the adiabatic approximation and build in memory or maybe even initial state dependence. For benchmarking and gaining insight in the construction of such new approximations the ability to construct the exact Hxc potential is very valuable. It allows us to see precisely what the approximate Hxc potential is missing and which features the exact potential possesses.

7.2 2D model system

We now consider the case of interacting and non-interacting systems in two dimensions. For this purpose we consider an analogous model system to the one-dimensional model system that we discussed above. We use the potential v_0 (to obtain n_0) and interaction w given by[21]

$$v_0(x,y) = -\cos\left(\tfrac{2\pi x}{L}\right) - \cos\left(\tfrac{2\pi y}{L}\right),$$
$$w(x_1,y_1,x_2,y_2) = \lambda\cos\left(\tfrac{2\pi(x_1-x_2)}{L}\right) + \lambda\cos\left(\tfrac{2\pi(y_1-y_2)}{L}\right),$$

and still take $L = 10$ (both sides) and $T = 20$ and again use periodic boundary conditions. We use these potentials to compute the initial state and construct a time-dependent density

$$n(xyt) = \frac{1}{4}[n_0(x - r(t), y) + n_0(x + r(t), y) + n_0(x, y - r(t)) + n_0(x, y + r(t))],$$

which describes a splitting in four fragments (instead of only two as in the one-dimensional case) and where we

[21] This interaction is a bit unusual since it depends not only on the distance between the particles and therefore is different for $x_1 - x_2 = \sqrt{2}$, $y_1 - y_2 = 0$ and $x_1 - x_2 = 1$, $y_1 - y_2 = 1$. However, the algorithm can also handle such a case.

have used the same $r(t)$ of equation (20) as in the 1D case. Finally we again use our algorithm to compute $v[\Psi_0, n]$.

Note that in all cases we only provide 6 snapshots of the potentials though the potentials change much more in time than this represents. In all cases $v(\mathbf{r}t_0) = v_0(\mathbf{r})$ since $\partial_t^2 n(\mathbf{r}t_0) = 0$ as explained in Section 7.1. As a side remark, we note that the given v_0 and w decouple in x and y components, such that the spatial parts of the eigenstates of the Hamiltonian are products of the individual eigenstates of the x and y components of this Hamiltonian (which are still interacting states if w is non-zero). While this simplifies the construction of the initial states it does not make the inversion easier as the time-dependent external potential is not separable in the two spatial coordinates. Finally, we note that we do not explicitly compute an exact Hxc potential here, as we find it more interesting to do a non-interacting inversion for a system of 10 electrons rather than for the two-electron case needed to compute the Hxc potential (where one could also use the analytic single electron formula).

7.2.1 One electron

In our first case, we simply consider a single electron and take as our initial state one of the (spatially identical) degenerate spin-up or spin-down ground states (whose spatial part is simply a product of identical functions of x and y each of which is the one-electron ground state of the one-dimensional system in Sect. 7.1.1). Snapshots of the resulting potential are provided in Figure 7. As expected for a single electron case the potential is periodic in time and the periodicity of the density caries over to the potential (as do the other symmetries of the densities). Moreover, at the final time we recover the original state as in the one-dimensional case.

7.2.2 Interacting electrons

In our second case we take the initial state to be the two-electron singlet ground state (whose spatial part is simply a product of identical functions of x_1, x_2 and y_1, y_2, that each are the 2-electron ground states of the system in Sect. 7.1.2). Snapshots of the resulting potential are provided in Figure 8. Just like in the one-dimensional case the potential is now no longer periodic in time and we no longer recover the original state at the final time.

7.2.3 Non-interacting electrons

We now consider the non-interacting 10-electron ground state for which the one-particle orbitals can be constructed as products of the orbitals of Figure 4. Snapshots of the resulting potential are provided in Figure 9. As expected the potential is again not periodic in time and we end up in a different final state at the final time.

7.3 3D model system

To demonstrate that the method also performs well in the practically relevant three-dimensional case let us finally consider a 3D system of 14 non-interacting electrons with

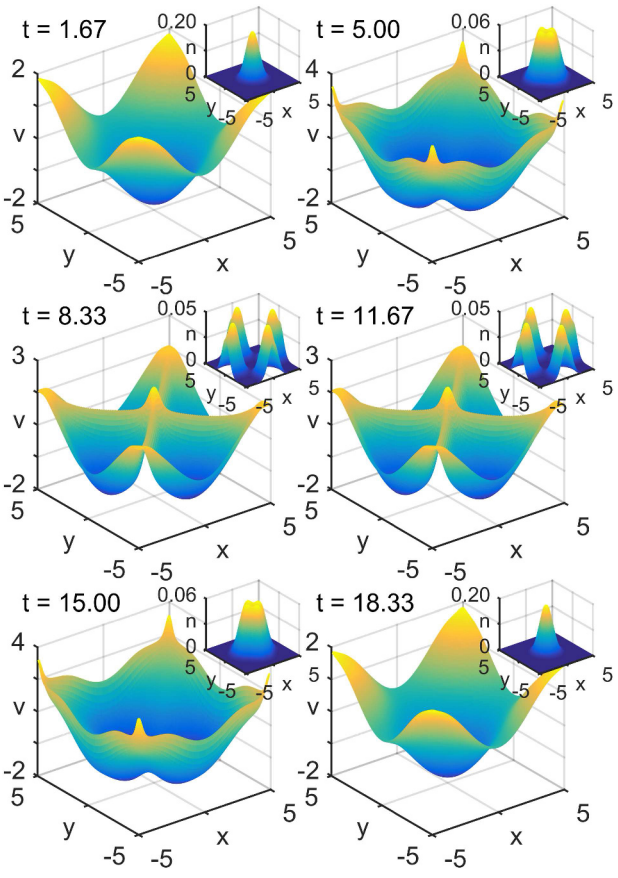

Fig. 7. Snapshots of the potential and density (insets) for a single 2D electron at different times.

an initial one-body potential

$$v_0(xyz) = -\cos\left(\tfrac{2\pi x}{L}\right) - \cos\left(\tfrac{2\pi y}{L}\right) - \cos\left(\tfrac{2\pi z}{L}\right).$$

This time we take $L = 5$ (all sides) and $T = 5$ to make the numerical data files smaller for convenience. We again use periodic boundary conditions and compute the ground state initial state and density n_0 as before. We then construct a time-dependent density

$$\begin{aligned} n(xyzt) = \frac{1}{6}[&n_0(x - r(t), y, z) + n_0(x + r(t), y, z) \\ &+ n_0(x, y - r(t), z) + n_0(x, y + r(t), z) \\ &+ n_0(x, y, z - r(t)) + n_0(x, y, z + r(t))], \end{aligned}$$

which describes a fragmentation of the initial density in six pieces in which $r(t)$ is again given by equation (20). Again $v(\mathbf{r}t_0) = v_0(\mathbf{r})$, and v_0 and w decouples in x, y and z making it easier to construct the initial state although it does not simplify the time-dependent inversion. Snapshots of the resulting potential are provided in Figure 10 for $z = 0$. The $x = 0$ and $y = 0$ planes would look the same, while for example the $z = 1$ plane would differ, but we restrict to show $z = 0$ as it is not our purpose to make a full analysis of the three dimensional dynamics. As expected the potential is again not periodic in time and we do not end

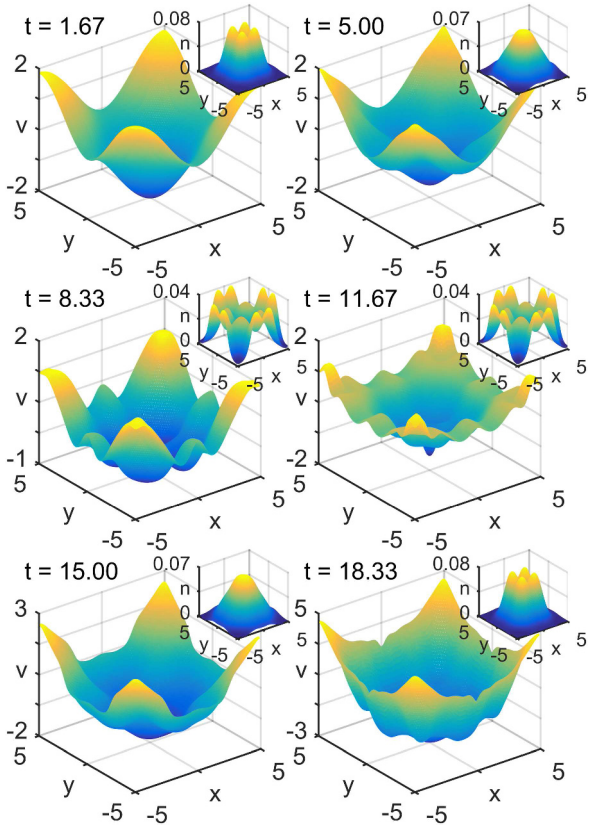

Fig. 8. Snapshots of the potential and density (insets) for 2 interacting 2D electrons at different times.

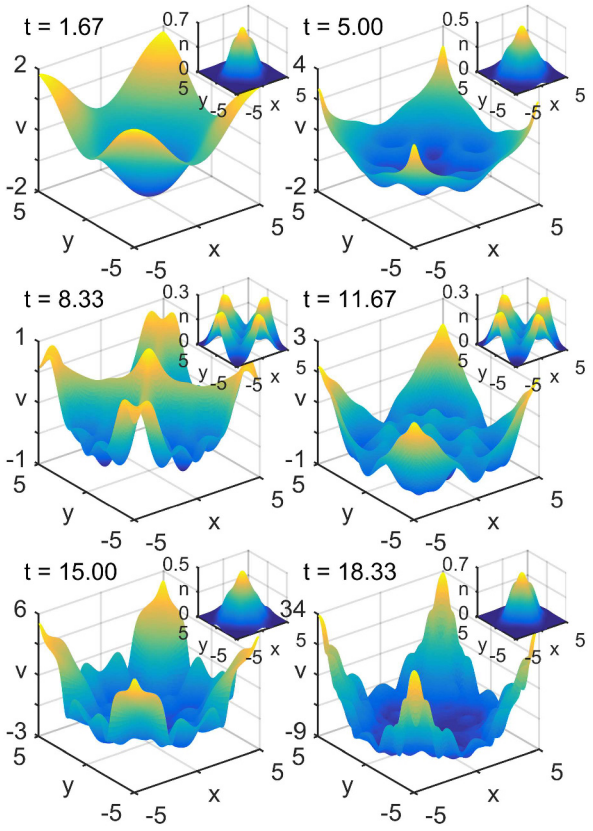

Fig. 9. Snapshots of the potential and density (insets) for 10 non-interacting 2D electrons at different times.

up back in the initial state. Furthermore the potential has relatively smooth features as we have many electrons in this case.

8 Conclusions and outlook

We have extended a recent inversion method [10] to 2D and 3D systems and presented the first inversions for 2D and 3D multi-electron systems. We discussed numerical details that make our method efficient, stable and precise even for large and rapid density changes (even by 34 orders of magnitude) irrespective of the initial state and the two-body interaction. We highlighted that we must treat the low density regions carefully, respect the boundary conditions, use a special time-stepping strategy, and explicitly enforce both the density and the current density to be accurate. We further showed that our method is very efficient, as it is only 5 to 10 times slower than the employed time-stepping method.

These results open up the possibility of a thorough investigation of density functionals in general, and the exact Hxc potential in particular, for realistic systems. This will be most valuable for the development of density functionals, and has for example already been used to study the possibility of using more advanced (multi-configurational) initial states in the KS approach [37]. In the context of strictly-correlated electron reference systems [40–42] it may help to open up new possibilities for the study of strongly correlated systems within a TDDFT framework. Furthermore, as already shown in, e.g., the context of cavity quantum electrodynamics [30], the method can be straightforwardly extended to different pairs of conjugate variables such as time-dependent current DFT or time-dependent dipole DFT [31].

The method is therefore a way to do rigorous tracking for the conjugate variables, and not only useful to study TDDFT. This quantum-control perspective will also be the subject of future work. Besides combining our method with many-body methods like TDDFT to predict approximate control fields for large systems we can also combine it with optimal-control theory. This combination allows us to search over a restricted set of densities instead of potentials as usual [43,44].

An important remaining challenge is to develop a complementary black box approach to optionally avoid the necessity to prescribe the density precisely in the low density regions. For example by designing artificial potentials for the low density regions that affect the high density regions minimally (although this is complicated). This would allow further approximations to be used, and thereby also to increase performance. Another challenge even then is to treat the moderately low density regions reasonable as discussed in Section 4. Many common approximations and the experience required to use

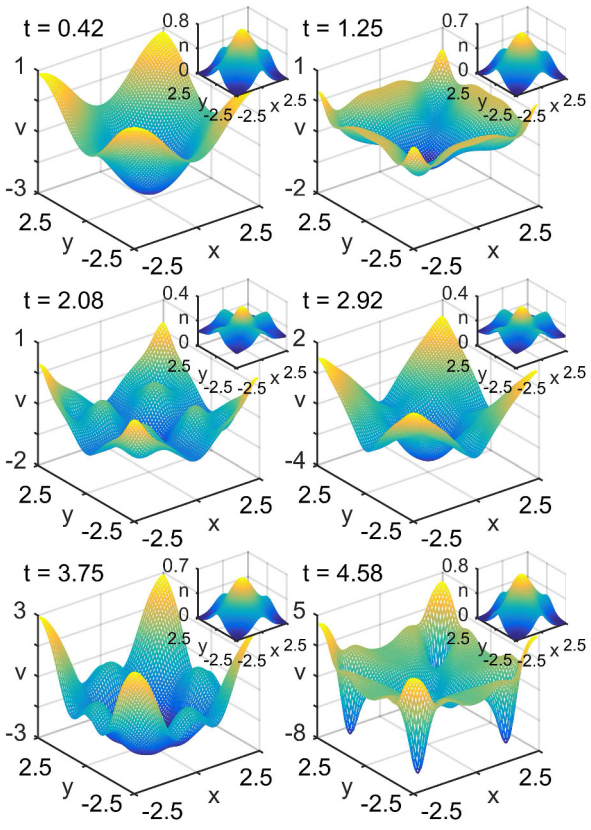

Fig. 10. Snapshots of the potential and density (insets) for 14 non-interacting 3D electrons at different times for the $z = 0$ plane.

them were not developed with these regions in mind. Additionally the dynamics in the low density regions is often much stronger than in the high density regions, as we for example saw in Figure 2a. There we obtained a much stronger potential in the low density region (sometimes the boundaries can also be the cause of strong dynamics). It is surprisingly easy to accidentally ask for a density change that requires a very strong potential causing our method to break down. For example, one may require a large density change or change of character in the wave function within only a few time-steps.

Open access funding provided by Max Planck Society. S.E.B. Nielsen and M. Ruggenthaler acknowledge financial support from the European Research Council (ERC-2015-AdG-694097) and European Union's H2020 programme under GA no. 676580 (NOMAD). We thank Dr. B. Ebner for computational support. We further thank Prof. Dr. A. Rubio for his support.

Author contribution statement

All authors actively participated in the theory development and the writing of the manuscript. S.E.B. Nielsen was responsible for the development of all the numerics, implementation and calculations.

Open Access This is an open access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Appendix A: Boundary conditions using analytic functions and finite difference formulas

The formulas employed in real-space codes, such as for derivatives or interpolation, usually presume that all functions can be represented by a convergent Taylor series, i.e., they assume analyticity of the functions. In this Appendix we therefore show how to work with analytic functions in a quantum mechanical context, and how to impose proper boundary conditions in this context. This allows us to, for example, derive consistent finite difference formulas using the analytic properties, and thereby to determine derivatives of wave functions, potentials or other quantities with high accuracy also at the boundary of the simulation box (necessary for inversions). It further allows us to make the boundary conditions a property of the physical quantities, e.g., the wave function or the density.[22] This is desirable, since then we can easily write generic computer routines, where we, for example, do not need to specify the boundary conditions for derivatives, but infer them from the objects. It further ensures an internal consistency when computing eigenstates, performing time-steps, and doing inversions, where the boundary conditions used for different routines always match. This is because (in a given system) densities, for example, always have the same "boundary conditions" property.

In this Appendix we will focus on the finite differences technique, but the ideas are equally applicable in other real space codes, and even to define a basis set. Without loss of generality, we further only consider a simple 3-point rule for a first-order derivative in one dimension,

$$f_i' = \tfrac{1}{2\Delta x}(f_{i+1} - f_{i-1}),$$

since the logics trivially generalises to higher-dimensions, higher-order derivatives or other stencils.[23] The only non-trivial aspect to apply such a derivative is the question what to do at the outermost points to get f'_{first} and f'_{last}. Consider, for instance, the sine and cosine functions of Figure A.1, which have a convergent Taylor expansion $f(x) = \sum_{n=0}^{\infty} c_n (x - \frac{L}{2})^n$ around $x = L/2$. If we take the boundary points at $\pm L/2$ then the derivatives at the last grid point in the interval are just

$$f'_{\text{last}} = \tfrac{1}{2\Delta x}(+f_{\text{last}} - f_{\text{last}-1}), \qquad \text{(A.1)}$$

[22] We point out that in standard quantum mechanics formulated on general square-integrable wave functions the boundary conditions are encoded in the set of allowed wave functions, i.e., the domain of the kinetic energy operator. Here we therefore determine the form of analytic wave functions in the domain of the kinetic energy operator.

[23] In practise, higher order stencils are generally preferable, especially in higher dimensions, since the higher precision allows one to use less grid-points. It thus saves work and memory in total, despite of the stencil itself being more involved.

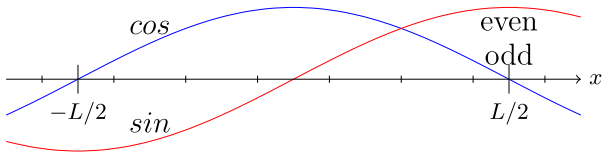

Fig. A.1. The sine and cosine functions. The grid points used for finite differencing are indicated by small vertical marks on the horizontal axis and the boundary points $\pm L/2$ are exactly halfway between two grid points.

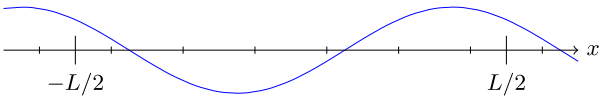

Fig. A.2. Example of a periodic function with grid points indicated by small vertical marks and the boundary points $\pm L/2$ exactly halfway between two grid points. For finite differencing at the last grid point before the boundary at $L/2$ we can use equation (A.3).

$$f'_{\text{last}} = \tfrac{1}{2\Delta x}(-f_{\text{last}} - f_{\text{last}-1}), \tag{A.2}$$

since the neighbouring value to the last value equals itself in the sine (even with respect to the boundary) case, and minus itself in the cosine (odd at the boundary) case (see Fig. A.1). The cosine (odd) case, being the ground-state of a box of length L with zero boundary conditions, already gives a hint as how to work with zero boundary conditions. But let us go through this in a little more detail.

A.1 Periodic boundary conditions

Let us consider the simpler periodic case first. Since all derivatives are continuous for analytic functions, the periodic boundary conditions for a function $f(x)$ on an interval $[a, b]$ imply that $\partial_x^n f(a) = \partial_x^n f(b)$ for all $n \in \mathbb{N}$. Indeed, since we know that all eigenfunctions of a Schrödinger problem with analytic potentials are analytic again [45], this implies that for periodic analytic potentials the eigenfunctions obey the above periodicity conditions. Further, if we perform a time-independent inversion for an analytic and periodic density the resulting potential is analytic and periodic as well. Also, if we propagate an analytic and periodic initial state with an analytic and periodic midpoint potential the resulting time-dependent wave function will stay analytic and periodic and we get the same answer for both time-step operators, i.e., equations (9) and (10).[24] Finally, given an analytic and periodic initial state and density a time-dependent inversion will lead to an analytic and periodic potential.[25] We therefore see that in the periodic case, by restriction to only periodic and analytic functions we gain a nice internal consistency of all the different components that appear for inversions (and beyond).

Now, if a quantity is analytic and periodic as in Figure A.2, the finite differencing at the last point becomes

$$f'_{\text{last}} = \tfrac{1}{2\Delta x}(f_{\text{first}} - f_{\text{last}-1}), \tag{A.3}$$

where we merely use the second last and the very first point of the interval. For a wave function to stay periodic and analytic also the applied potentials and interaction need to be analytic and periodic. If we, in the continuum case, prescribe a non-periodic or non-analytic potential or interaction the above finite differencing will essentially interpret it as a periodic and analytic approximation that depends on the grid-spacing. This usually works well as long as the violation of the analyticity and periodicity conditions is not too strong, e.g., the potential jumps discontinuously. The reason for this is that we can approximate any square-integrable function by an analytic function to any accuracy we want. Although the approximation smoothens the non-analyticities (allowing the approximation to work in the first place), it will still be close to a function that yields the right result. However, if we try to approximate functions that are not differentiable (already in the continuum problems can arise when acting with the kinetic-energy operator) then also the discretization becomes problematic.

A.2 Zero boundary conditions

For the zero-boundary situation (hard-wall boundary conditions) we let us be inspired by the example above for the cosine (odd) and sine (even) cases. To find the same internal consistency like in the periodic case, we consider wave functions that are odd across the boundary as they have a node at the boundary, and hence the derivative is given by equation (A.2). For the density (proportional to the modulus square of the wave function) we accordingly find that it is even across the boundary and the derivative is given by equation (A.1). The current density (being proportional to the wave function times its first derivative) is odd across the boundary. To keep the necessary even-odd distinction of the different physical quantities also the potential and interaction need to be even across the boundary.

To see why one should require these even-odd restrictions, let us first consider the case of free propagation with the Hamiltonian $\hat{H}(t) = \hat{T}$ and zero boundary conditions. For the wave function to stay zero at the boundary during the time propagation with equation (10) the wave function must clearly stay odd in time over the boundary. If it has any even term in its Taylor expansion around $x = \pm L/2$ some $\hat{T}^n \Psi$ will become non-zero at the boundary and so will the wave function after an infinitesimal time. The odd-ness condition is an important property as it is not sufficient for the initial state to merely be zero at the boundary if the zero-boundary condition is supposed to survive throughout the free propagation.

To ensure that an odd wave function stays zero at the boundary when propagating with some potential and interaction these are required to be even. In this case $\hat{H}$

[24] Although we do not have a strict proof for this, recent results in the context of the mathematical formulation of the Runge–Gross result [23,28] suggest that such a theorem should be possible.

[25] Again we have no strict proof, but under certain assumptions the analyticity of the potentials and wave functions are connected [20].

is even as well such that all $\hat{H}^n\Psi$ are odd and the wave function stays odd. If $\hat{H}$ also has odd components then even terms mix into $\hat{H}\Psi$ and, like in the case of free propagation, some higher powers $\hat{H}^n\Psi$ will become non-zero at the boundary.[26]

These results establish an internal consistency of odd-even relations. Wave functions, densities, potentials and interactions cannot be chosen independently but need to be consistent with each other. These even-odd relations are the ones we use in practice. Of course, it is in principle possible to go beyond these simple consistency relations by, e.g., using an initial state that is not odd but still zero at the boundaries. To then ensure that the wave function stays zero when propagating with equation (10) we need some specially designed potentials.[27] Whether a similar explicit characterization of internal consistency as in the even-odd case is possible also in this case and the two time-step operators equations (10) and (9) still agree is the subject of future research. However, since we can approximate any square-integrable function by corresponding odd (for wave functions) or even (for potentials and interactions) analytic functions, the even-odd analytic functions are general enough to cover almost all cases of physical interest (the subsequent numerical discretization limits these somewhat which still encourages interest in other than even-odd functions). For instance, the soft Coulomb potential in 1D can be easily handled but more pronounced violations of analyticity conditions, e.g., discontinuous jumps in the potentials, might lead (as discussed in the periodic case) to problems.

We finally want to make a few important remarks. The even and odd analytic functions are periodic on an interval of twice the length as shown in Figure A.3. When we ignore the even-odd consistency conditions the wave function tends to become noisy and inaccurate near the boundary. In standard time-propagation most observables will essentially be unaffected since the noise usually occurs

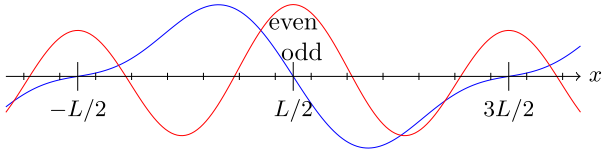

Fig. A.3. All analytic functions which are even or odd with respect to the boundaries at $\pm L/2$ are periodic on the double interval $[-L/2, 3L/2]$.

in the low density regions. However, if the density is not small near the boundaries, as in the case of a reflection, these small errors affect accuracy and numerical stability. Further, if all quantities already adhere to the even-odd conditions, one can mathematically also use zero boundary conditions on the wave function and no boundary conditions on the potentials, and it will keep the even-odd conditions. Numerically, this corresponds to using sided finite differences for the potential and sided finite differences with the zero value at the boundary build in for the wave function. Whether this is sufficient to ensure stability numerically for time-dependent inversion, or one has to explicitly enforce the even-odd conditions, remains an interesting open question since these sided finite differences are also natural to use in the time-independent case. Note that we place the boundary points in Figure A.1 halfway between the grid points instead of including them in the equidistant grid (which is common for the sided finite differences and sided finite differences with the zero value build in). This avoids to either store the zero values for the wave function at the boundary, complicating division, or to use a different number of grid points for the potential and wave function (since for the potential we need to store the values at the boundary to take, for example, derivatives thereof). In this way we can also use the same grid for periodic boundary conditions, without having a "boundary" point that is covered twice, making re-use of the code more elegant and easier.

Appendix B: Vanishing density regions

In this Appendix we provide the basic ideas that allow us to also treat vanishing density regions to a high precision. We stress that one can also treat most cases without these ideas, but they do allow for a broader class of problems. Note that all the propagation and inversion methods we present in the main text allow for vanishing densities, as long as one computes the initial state and prescribed density with sufficient precision. Therefore this Appendix is mainly concerned with the time-independent Schrödinger equation.

B.1 Vanishing density methods

The time-independent Schrödinger equation is usually solved iteratively like other large eigenvalue problems. The typical solvers take a number of trial vectors, $\{\psi_i\}$, and update them repeatedly until they all (or a desired amount thereof) converge (it is often faster to use more trial vectors than necessary). To obtain a high accuracy even for

[26] More strictly, if $v(x) = \sum_{n=0}^{\infty} v_n x^n$ and $\Psi(x) = \sum_{n=0}^{\infty} \Psi_n x^n$ are Taylor expansions at the boundary and $v(x)$ has a lowest odd term v_{2n+1} and $\Psi(x)$ is odd, then $\hat{H}^{n+1}\Psi(0) = \frac{(2n)!}{(-2)^n} v_{2n-1}\Psi_1$. This term arises from $\hat{T}^n\hat{V}\Psi(0)$ and is the only term involving an odd v_n that survives, i.e., sufficient derivatives are needed to end at zero order (x^0), which is the only order not to vanish at the boundary (as for all other $x^n = 0$). All the terms involving only even v_n vanish at the boundary as they are odd. For $\hat{H}^{n+1}\Psi(0)$ to vanish we therefore need $v_{2n-1} = 0$ and so all odd v_n must indeed vanish. Note even if $\Psi_1 = 0$ one can use $\hat{H}^{n+2}$ instead to obtain a similar result with Ψ_3, and similarly one can find higher order expression also, when also Ψ_3 or higher odd powers vanish.

[27] If an initial state is not odd, we of course still need $\hat{H}^n\Psi$ to vanish at the boundary. Let us consider one electron in one dimension with Hamiltonian $\hat{H} = -\frac{1}{2}\partial_x^2 + v(x)$. By Taylor expansion at the boundary, $v(x) = \sum_{n=0}^{\infty} v_n x^n$ and $\Psi(x) = \sum_{n=0}^{\infty} \Psi_n x^n$ this leads to the conditions $\Psi(0) = \Psi_0 = 0$, $\hat{H}\Psi(0) = -\Psi_2 = 0$, $\hat{H}^2\Psi(0) = 6\Psi_4 - v_1\Psi_1 = 0$, $\hat{H}^3\Psi(0) = -90\Psi_6 + 6v_3\Psi_1 + 9v_1\Psi_3 = 0$, $\hat{H}^4\Psi(0) = 2520\Psi_8 - 90v_5\Psi_1 - 108v_3\Psi_3 - 36v_2\Psi_4 - 180v_1\Psi_5 = 0$, and so on, where we have used the previous relations to simplify the latter. These equations tell us which potentials are allowed for a given wave function and vice versa, but they are not very practical and we do not know whether they provide separate conditions on $v(x)$ and $\Psi(x)$ (especially for higher orders).

very low values like 10^{-27} one needs to use an extremely strict convergence criterion and a method able to reach this extreme level of precision.

As convergence criterion we advice to require that $\hat{H}|\psi_i\rangle = E_i|\psi_i\rangle$ pointwise to a high relative precision where E_i is computed by $\langle\psi_i|\hat{H}|\psi_i\rangle$. To enforce this we require that $|(\hat{H} - E_i)\psi_i(x)|/|\psi_i(x)| \leq \delta$ for all x (or for all space and spin coordinates in general), with for example $\delta = 10^{-8}$. If fulfilled both eigenstates and energies must clearly be correct to a high relative precision everywhere, also for very small values of the order of 10^{-27} and even smaller when it is relevant. This is a very natural criterion to use (for the high density region one may require $|(\hat{H} - E_i)\psi_i(x)| \leq \delta$ instead). For excited states it is generally advantageous to only require the criterion for a fraction of the grid-points, such as 0.96, because, due to nodes and division by ψ_i in the criterion, there are often a few points where it is hard to reach as high precision as everywhere else.

For the solver we recommend polynomial filtering [46], for which imaginary Euler time-stepping is a very simple option. In this approach one uses the update $\psi_{i,k+1} = \psi_{i,k} - \Delta t\hat{H}\psi_{i,k}$ followed by a re-normalization. Given multiple trial vectors, a few re-orthogonalizations will also be needed to ensure that each trial vector gives a different of the lowest eigenstates rather than all the lowest. We recommend this method as a first method as it practically always works. However, it is also very slow, and therefore we also recommend to build a more advanced polynomial filter for a large performance boost. For high performance the filters need some fine-tuning to specific cases, especially when computing many eigenstates, but are quite special in how they can achieve almost any accuracy for almost any system and number of eigenstates. For a good filter to compute ground states or a number of the lowest eigenstates we refer to [47].

Note that usual variational methods will generally fail. When minimizing the energy $E_0 =< \psi_0|\hat{H}|\psi_0 >$ the vanishing density regions contribute very little and therefore when E_0 is limited to double precision (about 16 digits) these small contributions are not taken into account and the noise in the tails of the eigenstates persists. Many standard methods, including the very popular Lanczos method for computing eigenstates (which most library solvers use) and many optimization methods for time-independent inversion [15], share similar issues. Unless higher floating point precision is used these methods are not accurate enough in the vanishing density context. However, we recommend the Lanczos method in cases where the density is not vanishing as it is easy to implement and very effective.

Another situation to avoid in the vanishing density context is to take a sum of large values to obtain a very small value. It is not possible to obtain a value of say 10^{-20} precisely, if we perform a sum that involves a value of say 1 to get it, as the value of 1 is only precise to about 10^{-16}. Even when not working with vanishing densities we generally recommend to use a numerically more precise way to sum values known as a Kahan sum whenever computing inner products or expectation values. This prevents the individual truncation errors from accumulating so the sum stays precise even when it involves billions of terms. For vanishing densities, one should be very careful if not only using the values at relatively local grid points to for example compute the derivatives of a wave function, as the global values may be many orders of magnitude larger. Furthermore one should, for example, not take Fourier transforms to momentum space and back. While mathematically this should act as an identity transform, numerically this is performed by taking a sum of the values at all grid points which may again involve significantly larger values than the final value. Such considerations strongly limit the use of global information in the vanishing density context, which is one of the largest drawbacks of having vanishing densities as it limits the choice of methods. However, all of our methods are designed to easily avoid these issues if one uses local representations such as finite difference techniques.

B.2 Alternative representations

Of course, an alternative to using special methods that allow us to go to extremely low values, is to simply change to a representation where the values never become very small and thereby circumvent the issue. Indeed, there are few good reasons not to use a basis representation and avoid such very low values in the time-independent case. Using basis functions is also a nice way to deal with the boundary conditions, or to work on an infinite domain. However, note that with basis functions it can be hard to accurately represent the low-density regions. This is because the coefficients of the basis functions (which have a fixed form in the low-density region) are mainly fixed by the high-density region. It is therefore essential that the basis functions decay in a physical manner in the low density regions to not become too unphysical (which is problematic for inversions). In fact the basis functions should even be rather physical in the high density regions, or it can lead to artificially oscillating potentials if one is not cautious [48,49].

In the time-dependent case it may also be easy in given cases to avoid very low values by changing representation, although this is much more involved for a general purpose solver. For example, the wave function may change by orders of magnitude over time at some points which will complicate the situation considerably. In particular, we expect that it will be very hard to obtain the tails of the wave function sufficiently accurately, and errors may quickly grow and propagate into the relevant region as we step through time.

References

1. U. von Barth, Phys. Scr. **T109**, 9 (2004)
2. C.A. Ullrich, *Time-Dependent Density-Functional Theory: Concepts and Applications* (Oxford University Press, Oxford, 2012)
3. M.A.L. Marques, C.A. Ullrich, F. Nogueira, A. Rubio, K. Burke, E.K.U. Gross, *Time-Dependent Density Functional Theory* (Springer, Heidelberg, 2012)
4. N.T. Maitra, K. Burke, Phys. Rev. A **63**, 042501 (2001)

5. N.T. Maitra, K. Burke, C. Woodward, Phys. Rev. Lett. **89**, 023002 (2002)
6. M. Ruggenthaler, S.E.B. Nielsen, R. van Leeuwen, Phys. Rev. A **88**, 022512 (2013)
7. P. Elliott, J.I. Fuks, A. Rubio, N.T. Maitra, Phys. Rev. Lett. **109**, 266404 (2012)
8. M.J.P. Hodgson, J.D. Ramsden, J.B.J. Chapman, P. Lillystone, R.W. Godby, Phys. Rev. B **88**, 241102 (2013)
9. M.J.P. Hodgson, J.D. Ramsden, R.W. Godby, Phys. Rev. B **93**, 155146 (2016)
10. S.E.B. Nielsen, M. Ruggenthaler, R. van Leeuwen, Europhys. Lett. **101**, 33001 (2013)
11. C. Verdozzi, Phys. Rev. Lett. **101**, 166401 (2008)
12. J.D. Ramsden, R.W. Godby, Phys. Rev. Lett. **109**, 036402 (2012)
13. P. Schmitteckert, M. Dzierzawa, P. Schwab, Phys. Chem. Chem. Phys. **15**, 5477 (2013)
14. J.D. Whitfield, arXiv: 1503.00248v1 (2015)
15. D.S. Jensen, A. Wasserman, Phys. Chem. Chem. Phys. **18**, 21079 (2016)
16. P. Gross, H. Singh, H. Rabitz, K. Mease, G.M. Huang, Phys. Rev. A **47**, 4593 (1993)
17. W. Zhu, H. Rabitz, J. Chem. Phys. **119**, 3619 (2003)
18. M. Ruggenthaler, R. van Leeuwen, Europhys. Lett. **95**, 13001 (2011)
19. M. Ruggenthaler, K.J.H. Giesbertz, M. Penz, R. van Leeuwen, Phys. Rev. A **85**, 052504 (2012)
20. M. Ruggenthaler, M. Penz, R. van Leeuwen, J. Phys.: Condens. Matter **27**, 203202 (2015)
21. V. Peuckert, J. Phys. C: Solid State Phys. **11**, 4945 (1978)
22. E. Runge, E.K.U. Gross, Phys. Rev. Lett. **52**, 997 (1984)
23. S. Fournais, J. Lampart, M. Lewin, T.Ø. Sørensen, Phys. Rev. A **93**, 062510 (2016)
24. C. Leforestier et al., J. Comput. Phys. **94**, 59 (1991)
25. R. van Leeuwen, E.J. Baerends, Phys. Rev. A **49**, 2421 (1994)
26. D.S. Jensen, A. Wasserman, Int. J. Quantum Chem. **118**, e25425 (2017)
27. W. Kohn, L.J. Sham, Phys. Rev. **140**, A1133 (1965)
28. M. Penz, arXiv: 1610.05552v1 (2016)
29. A. Castro, M.A.L. Marques, A. Rubio, J. Chem. Phys. **121**, 3425 (2004)
30. J. Flick, M. Ruggenthaler, H. Appel, A. Rubio, Proc. Natl. Acad. Sci. U.S.A. **112**, 15285 (2015)
31. M. Ruggenthaler, M. Penz, D. Bauer, Phys. Rev. A **81**, 062108 (2010)
32. M. Penz, M. Ruggenthaler, J. Chem. Phys. **142**, 124113 (2015)
33. M. Penz, arXiv: 1801.03361 (2018)
34. J.I. Fuks, L. Lacombe, S.E.B. Nielsen, N.T. Maitra, arXiv: 1806.10267 (2018)
35. W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, Cambridge, 2007)
36. I. D'Amico, G. Vignale, Phys. Rev. B **59**, 7876 (1999)
37. J.I. Fuks, S.E.B. Nielsen, M. Ruggenthaler, N.T. Maitra, Phys. Chem. Chem. Phys. **18**, 20976 (2016)
38. Y. Suzuki, L. Lacombe, K. Watanabe, N.T. Maitra, Phys. Rev. Lett. **119**, 263401 (2017)
39. L. Lacombe, Y. Suzuki, K. Watanabe, N.T. Maitra, Eur. Phys. J. B **91**, 96 (2018)
40. M. Seidl, Phys. Rev. A **60**, 4387 (1999)
41. M. Seidl, P. Gori-Giorgi, A. Savin, Phys. Rev. A **75**, 042511 (2007)
42. F. Malet, A. Mirtschink, K.J.H. Giesbertz, L.O. Wagner, P. Gori-Giorgi, Phys. Chem. Chem. Phys. **16**, 14551 (2014)
43. W. Zhu, J. Botina, H. Rabitz, J. Chem. Phys. **108**, 1953 (1998)
44. I. Serban, J. Werschnik, E.K.U. Gross, Phys. Rev. A **71**, 053810 (2005)
45. S. Fournais, M. Hoffmann-Ostenhof, T. Hoffmann-Ostenhof, T.Ø. Sørensen, Ark. Mat. **42**, 87 (2004)
46. A. Pieper, M. Kreutzer, A. Alvermann, M. Galgon, H. Fehske, G. Hager, B. Lang, G. Wellein, J. Comput. Phys. **325**, 226 (2016)
47. Y. Zhou, Y. Saad, SIAM J. Matrix Anal. Appl. **29**, 954 (2007)
48. P.R.T. Schipper, O.V. Gritsenko, E.J. Baerends, Theor. Chem. Acc. **98**, 16 (1997)
49. I.G. Ryabinkin, S.V. Kohut, V.N. Staroverov, Phys. Rev. Lett. **115**, 083001 (2015)