Pseudopotential for the 2D contact interaction

T.M. Whitehead,[1] L.M. Schonenberg,[1] N. Kongsuwan,[2] R.J. Needs,[1] and G.J. Conduit[1]

[1] *Cavendish Laboratory, J.J. Thomson Avenue, Cambridge, CB3 0HE, United Kingdom*
[2] *Blackett Laboratory, Prince Consort Road, London, SW7 2AZ, United Kingdom*
(Dated: March 16, 2016)

We propose a smooth pseudopotential for the contact interaction acting between ultracold atoms confined to two dimensions. The pseudopotential reproduces the scattering properties of the repulsive contact interaction up to 200 times more accurately than a hard disk potential, and in the attractive branch gives a 10-fold improvement in accuracy over the square well potential. Furthermore, the new potential enables diffusion Monte Carlo simulations of the ultracold gas to be run 15 times quicker than was previously possible.

I. INTRODUCTION

Many collective quantum phenomena emerge from reduced dimensionality, including the quantum Hall effect [1], high-temperature superconductivity [2], quantum magnetism [3], and topological insulators [4]. Consequently, two-dimensional (2D) systems have recently attracted a great deal of attention [5–11]. 2D systems may now be realized, for example, at the interface between two solids [12], or in an ultracold atomic gas in an anisotropic optical trap, with one dimension tightly confined relative to the other two [13–15]. This coincidence of novel many-body phenomena with accurate experimental realizations makes 2D systems attractive for numerical investigation.

Ultracold atoms provide a clean model Hamiltonian with a tunable interaction strength, and their study has delivered new insights into many-body quantum physics [16–19]. The resonant Feshbach interaction [20] between ultracold atoms is usually modeled by a contact potential [21–29]. Despite its widespread usage, the contact interaction causes sampling problems in numerical simulations due to its infinitesimally short range and divergence at coalescence. It also harbors a bound state, complicating the use of ground state methods for examining repulsive scattering between particles. These difficulties are conventionally circumvented by replacing the contact potential by, for example, a hard disk potential, which we show leads to inaccurate scattering properties. This problem has recently been resolved in three dimensions by Bugnion *et al.* with the development of a smooth pseudopotential [30] that results in a hundred-fold increase in the accuracy of the scattering properties. The smoothness of the new pseudopotential also radically speeds up numerical calculations [31, 32]. Here we follow that prescription to develop a pseudopotential that improves the modeling of 2D quantum gases with a contact interaction.

In Section II we discuss two particles interacting via the 2D contact potential. In Section III we derive the pseudopotential, and in Section IV demonstrate its accuracy in an inhomogeneous two-body system. In Section V we examine the pseudopotential's advantages over other methods in a homogeneous many-body system, before discussing potential future applications of the pseudopotential in Section VI.

II. ANALYTICAL RESULTS

In order to develop a pseudopotential for use in many-body simulations, it is essential to first understand the behavior of the two-particle system. Here we analyze an isolated two-body system of distinguishable fermions, starting with non-interacting particles and then adding a short-ranged interaction potential, which not only allows us to find solutions for the contact interaction, but also serves as a platform from which to propose a pseudopotential. Atomic units ($\hbar = m = 1$) are used throughout, and anticipating that we will study many-body systems, we measure energies in units of the Fermi energy E_{F} and lengths in units of the Fermi length k_{F}^{-1}.

A. Short-ranged two-particle interactions

We consider two equal-mass, distinguishable fermions in a vacuum. In their center-of-mass frame, the Schrödinger Equation for particles interacting via a potential $V(r)$ is given by

$$-\nabla^2 \psi(r,\theta) + V(r)\psi(r,\theta) = E\psi(r,\theta), \qquad (1)$$

where E is the energy in the center-of-mass frame.

The analytic solution to Equation (1) for non-interacting particles ($V(r) = 0$) in a vacuum takes the form

$$\psi_\ell(r,\theta) = R_\ell(r)\Theta_\ell(\theta)$$

with

$$\Theta_\ell(\theta) = \frac{1}{\sqrt{2\pi}}e^{i\ell\theta},$$
$$R_\ell(r) = \mathcal{A}(k)J_\ell(kr) + \mathcal{B}(k)Y_\ell(kr),$$

where $k = \sqrt{E}$ is the wavevector in the center-of-mass frame, ℓ is angular momentum projected onto the normal to the 2D plane, and $\mathcal{A}(k)$ and $\mathcal{B}(k)$ are constants set by

the boundary conditions. $J_\ell(kr)$ and $Y_\ell(kr)$ are Bessel functions of the first and second kinds, respectively.

If we take the potential $V(r)$ to be short ranged and cylindrically symmetric, for distinguishable fermions the only effect of the potential is in the $\ell = 0$ channel. The wavefunction beyond the interaction range, where $V(r) = 0$, then takes the same form as in the non-interacting case,

$$\psi_0(r) \propto \mathcal{A}(k)J_0(kr) + \mathcal{B}(k)Y_0(kr). \qquad (2)$$

There are two branches of solutions, scattering states for $E > 0$ and bound states for $E < 0$.

1. Scattering states (E > 0)

For two-particle scattering with positive E, the non-interacting wavefunction given by Equation (2) with $k = \sqrt{E}$ can be written at large separations in the oscillatory form

$$\psi_s(r) \propto \frac{\sin\left(kr + \frac{\pi}{4} + \delta(k)\right)}{\sqrt{kr}}, \qquad (3)$$

where the scattering phase shift $\delta(k)$, given by

$$\cot\delta = -\mathcal{A}(k)/\mathcal{B}(k), \qquad (4)$$

describes the large radius behavior of the wavefunction and captures the full impact of the scattering interaction.

2. Bound states (E < 0)

Two particles with $E < 0$ are in a bound state in which they remain in close proximity if no external force is applied. The bound state wavefunction has the form

$$\psi_b(r) \propto \mathcal{A}(\kappa)J_0(i\kappa r) + \mathcal{B}(\kappa)Y_0(i\kappa r),$$

where $\kappa = \sqrt{-E}$. For the wavefunction to be normalizable $\mathcal{A}$ and $\mathcal{B}$ must satisfy $\mathcal{B}(\kappa)/\mathcal{A}(\kappa) = i$, and therefore the wavefunction

$$\psi_b(r) \propto J_0(i\kappa r) + iY_0(i\kappa r)$$
$$\propto H_0^{(1)}(i\kappa r), \qquad (5)$$

where $H_0^{(1)}(x) = J_0(x) + iY_0(x)$ is the Hankel function of the first kind. Note that $H_0^{(1)}(i\kappa r) \to e^{-\kappa r}/\sqrt{\kappa r}$ as $\kappa r \to \infty$, with the expected exponential decay of a bound state.

B. 2D contact interaction

We now apply these results for short-ranged 2D interactions to the 2D contact interaction $V^{\mathrm{cont}}(r)$. In a
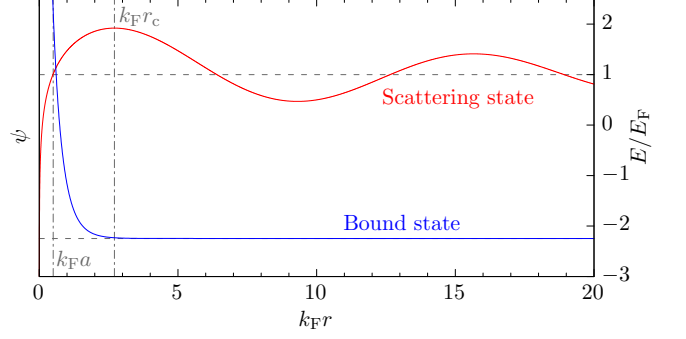

FIG. 1. (Color online) The bound (blue) and scattering (red) state wavefunctions of the contact interaction with $k_F a = 1/2$. The wavefunctions are offset by their energies, $E_s = E_F$ for the scattering state and $E_b = -\frac{4}{a^2}\exp(-2\gamma)$ for the bound state. The radius r_c gives the position of the first antinode in the scattering state, which is used as the cutoff radius for the pseudopotentials.

fermionic system this zero-ranged interaction acts only between distinguishable particles, with the interaction strength described by a scattering length a. We can capture the full effect of the interaction by imposing the boundary condition [33, 34]

$$\left(r\frac{d}{dr} - \frac{1}{\ln(r/a)}\right)\psi(r) = 0 \qquad (6)$$

at $r = 0$ and then at $r > 0$ use the non-interacting solution Equation (2), which gives

$$\psi_0^{\mathrm{cont}}(r) \propto J_0(kr) - \frac{\pi}{2[\gamma + \ln(ka/2)]}Y_0(kr), \qquad (7)$$

where $\gamma \approx 0.577$ is Euler's constant.

For $E > 0$ the scattering phase shift is evaluated using Equation (4) as

$$\cot\delta^{\mathrm{cont}}(k) = \frac{2}{\pi}\left[\gamma + \ln(ka/2)\right]. \qquad (8)$$

The pseudopotential must be able to reproduce this phase shift as a function of scattering wavevector.

For $E < 0$ the bound state wavefunction is given by Equation (5). The corresponding bound state energy can be found from the condition $\mathcal{B}/\mathcal{A} = i$ and Equation (7) as

$$E_b = -\frac{4}{a^2}e^{-2\gamma}.$$

Examples of a scattering and bound state wavefunction are shown in Fig. 1. At large radii the scattering state wavefunction takes the form of a wave, with the first node occurring at $r = a$ in the $k \to 0$ limit, whilst the bound state wavefunction decays exponentially. Both wavefunctions diverge at particle coalescence, which makes them difficult to sample in numerical methods. This motivates us to develop a smooth pseudopotential for the contact interaction, which will give rise to a wavefunction that is easier to sample numerically.

III.   DERIVATION OF THE PSEUDOPOTENTIALS

To develop smooth pseudopotentials for the contact interaction we continue to investigate the two-particle system in a vacuum, where the particles are distinguishable fermions and an analytical solution exists. We first focus on scattering states where, after reviewing the hard and soft disk potentials that are commonly used in ultracold atomic gas calculations, we construct a pseudopotential using the method proposed by Troullier and Martins (TM) [35]. This method was originally developed for making pseudopotentials for electron-ion interactions, but has been successfully applied to other systems of interacting particles [30, 32]. Next we construct another, "ultratransferable", pseudopotential (UTP) following the method in Ref. [30]. We then compare the accuracy of the TM and UTP pseudopotentials with that of the hard and soft disk potentials. Finally, we develop pseudopotentials for bound states. We have made the software used to generate all the pseudopotentials in this work available online [36].

A.   Pseudopotentials for scattering states

1.   Hard disk potential

Here we briefly review the hard disk potential that is currently used in many numerical studies of the contact interaction [11, 21, 22]. The interaction potential has the form

$$V^{\mathrm{HD}}(r) = \begin{cases} \infty, & r \le R, \\ 0, & r > R, \end{cases}$$

where R is the radius of the potential. Solving the Schrödinger Equation (1) with a boundary condition $\psi_0^{\mathrm{HD}}(R) = 0$, the wavefunction is given by

$$\psi_0^{\mathrm{HD}}(r) \propto \begin{cases} 0, & r \le R, \\ -Y_0(kR)J_0(kr) + J_0(kR)Y_0(kr), & r > R, \end{cases}$$

and the scattering phase shift defined by Equation (4) is

$$\begin{aligned} \cot \delta^{\mathrm{HD}}(k) &= \frac{Y_0(kR)}{J_0(kR)} \\ &= \frac{2}{\pi}[\gamma + \ln(kR/2)] + \frac{(kR)^2}{2\pi} + \mathcal{O}\big((kR)^4\big). \quad (9) \end{aligned}$$

By setting the hard disk radius R equal to the scattering length a, the low energy scattering phase shift from the hard disk has the same form as the phase shift from the contact potential in Equation (8). A hard disk potential with $R = a$ then gives the phase shift for the contact interaction with an error of order $\mathcal{O}\big((ka)^2\big)$, delivering the correct scattering properties only in the $k \to 0$ limit. An example of a hard disk potential is shown in Fig. 2.

2.   Soft disk potential

To reduce the error in the scattering phase shift at finite k from that found with the hard disk, we may instead use a soft disk potential [37]

$$V^{\mathrm{SD}}(r) = \begin{cases} U, & r \le R, \\ 0, & r > R, \end{cases}$$

where U is the height of the soft disk potential. The extra degree of freedom in this potential relative to the hard disk allows it to remove the error in the scattering phase shift in Equation (9) of $(kR)^2/2\pi$, and so describe the scattering correct to $\mathcal{O}\big((ka)^4\big)$. We solve the Schrödinger Equation using this potential separately in the regions $r < R$ and $r > R$, enforcing continuity of the wavefunction and its derivative at $r = R$, and expanding the scattering phase shift Equation (4) to second order around $k = 0$. Setting the first term equal to the contact potential scattering phase shift of Equation (8) relates R and the scattering length a via

$$R = a_{\mathrm{SD}} = a \exp\left(\frac{I_0(\chi)}{\chi I_1(\chi)}\right),$$

where $I_\ell(\chi)$ is the modified Bessel function of the first kind, and the factor $\chi^2 = UR^2 \approx 2.67$ is obtained by setting the second order term in the phase shift expansion to zero. This uniquely specifies a soft disk potential for a given a, whose scattering properties are correct up to order $\mathcal{O}\big((ka)^4\big)$. An example of a soft disk potential is shown in Fig. 2. It has a larger radius R than the hard disk potential but a lower height U, with the width tending to zero and the height to infinity as the scattering length goes to zero.

3.   Troullier–Martins pseudopotential

The previous subsections showed that the hard and soft disk potentials give accurate scattering properties only in the limit of $k \to 0$. However a Fermi gas contains all the scattering wavevectors in the range $0 < k \le k_{\mathrm{F}}$, and so the hard and soft disk potentials will give rise to inaccurate results. To demonstrate how the accuracy may be improved at finite k, we develop pseudopotentials using the TM formalism [30, 35]. This formalism produces scattering state pseudopotentials that

1. should reproduce the phase shift of the contact potential accurately for all scattering wavevectors in the Fermi sea;

2. are smooth everywhere, which accelerates numerical calculations; and

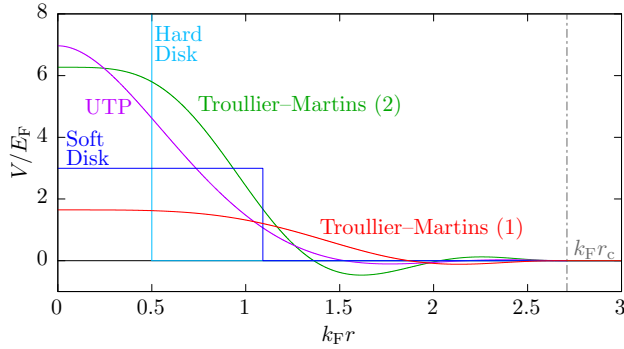3. for repulsive interactions do not support a bound state.

FIG. 2. (Color online) Scattering state pseudopotentials $V(r)$ for the contact potential with $k_F a = 1/2$. The pseudopotential for the hard disc with radius a is shown in cyan, the soft disk with radius a_{SD} in blue, the TM pseudopotentials in red and green, and the UTP in magenta. The pseudopotential cutoff radius r_c is the same as in Fig. 1.

FIG. 3. (Color online) The error in the scattering phase shift $|\delta^{pseudo}(k) - \delta^{cont}(k)|$ as a function of scattering wavevector for the different pseudopotentials at $k_F a = 1/2$. The error from the hard disk is shown in cyan, the error from the soft disk in blue, the errors from the two TM pseudopotentials in red and green, and the error from the UTP in magenta.

This formalism requires two prescribed parameters, namely the calibration wavevector k_c at which the resulting pseudo-wavefunction has identical scattering properties to the contact potential, and a cutoff radius r_c at which the pseudopotential smoothly becomes zero.

The calibration wavevector must be chosen for each system. For example, in a superfluid we might choose $k_c = k_F$, as that is where the most important physics of Cooper pair formation occurs. For a fermionic gas we choose $k_c = k_F/2$, which minimizes the average phase shift error [32].

By choosing the cutoff radius to be larger than the radius of the first node in the analytic wavefunction, which is at $r \approx a$ in Fig. 1, we ensure that the pseudo-wavefunction does not contain the innermost node that corresponds to the bound state of the contact interaction [30]. In order to avoid unnecessarily removing scattering states from the potential the cutoff radius must also be smaller than the radius of the second node, and so we choose the cutoff radius to be at the first antinode of the wavefunction with $k = k_c$, shown in Fig. 1.

The TM pseudopotential takes the form

$$V^{TM}(r) = \begin{cases} k_c^2 + p'' + p'^2 + \frac{p'}{r}, & r \leq r_c, \\ 0, & r > r_c, \end{cases}$$

where the polynomial $p(r) = \sum_{i=0}^{6} c_i r^{2i}$, and primes indicate derivatives with respect to r. The coefficients $\{c_i\}$ are determined by a set of constraints on the pseudopotential and pseudo-wavefunction, whose form is $\psi(r) = \exp[p(r)]$: that the pseudo-wavefunction is smooth up to the fourth derivative at r_c; that the pseudopotential has zero curvature at the origin; and that the norm of the pseudo-wavefunction within r_c equals that of the wavefunction from the real contact potential [30, 32, 35]. This gives rise to a set of coupled equations for the $\{c_i\}$ of which one is quadratic and the others linear: there are therefore two separate branches of
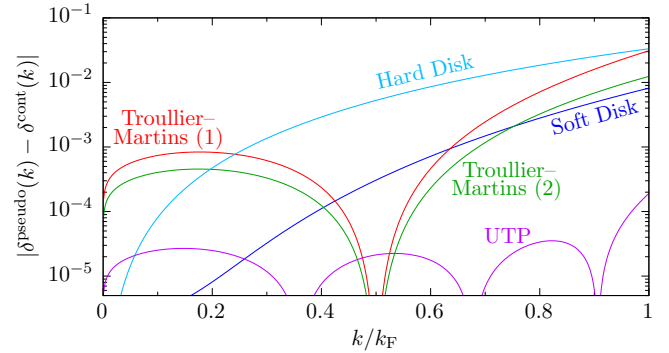
solutions, which give rise to two separate TM pseudopotentials.

In Fig. 2 we compare all of the discussed pseudopotentials for the contact interaction, with $k_F a = 1/2$. The TM pseudopotentials, being everywhere smooth and finite, are easier to work with numerically than the hard and soft disks, and they do not introduce discontinuities in the first derivative of the wavefunction. The potential labeled Troullier–Martins (1) in Fig. 2 is smaller than Troullier–Martins (2) at particle coalescence, but larger at further separations to give similar average scattering.

4. Errors in scattering phase shift

The quality of a pseudopotential for scattering states may be determined by how accurately it reproduces the phase shift of the contact potential. All information on the difference between the pseudopotential and contact potential can be obtained from the wavefunction just beyond the edge of the pseudopotential. We match the analytical pseudo-wavefunction that solves Equation (1), ψ, and its first derivative to the non-interacting solution Equation (2) at a radius R_e beyond the radius of the pseudopotential. This leads to an expression for the scattering phase shift

$$\cot \delta(k) = \frac{\frac{\psi'(R_e)}{\psi(R_e)} Y_0(k R_e) + k Y_1(k R_e)}{\frac{\psi'(R_e)}{\psi(R_e)} J_0(k R_e) + k J_1(k R_e)}.$$

We calculate the difference in the phase shift between the contact interaction and pseudopotentials, showing the error in the calculated phase shifts from using the pseudopotentials $|\delta^{pseudo}(k) - \delta^{cont}(k)|$ in Fig. 3, with $R_e = r_c$. The hard and soft disk potentials are exact in the limit of $k \to 0$, but deviate away from that point, with the soft disk performing better than the hard disk.

The TM (1) pseudopotential is on average around twice as accurate as the hard disk potential, with the TM (2) pseudopotential being around twice as accurate again, and the soft disk being another 1.3 times more accurate. Both TM pseudopotentials capture the scattering behavior perfectly at $k_c = k_F/2$ but deviate at all other scattering wavevectors, which is a consequence of the norm-conserving condition on the pseudo-wavefunctions. To further improve the accuracy of the pseudopotentials, a natural extension to the formalism is to construct a pseudopotential that minimizes this deviation in the phase shift over all wavevectors $k \leq k_F$. We propose such a pseudopotential here, referring to it as an "ultratransferable pseudopotential" (UTP) [30].

5. Ultratransferable pseudopotential

Similarly to the TM pseudopotential, the UTP takes a polynomial form within a cutoff radius r_c,

$$
V^{\mathrm{UTP}}(r) = \begin{cases} \left(1 - \frac{r}{r_c}\right)^2 \left[u_1\left(1 + \frac{2r}{r_c}\right) + \sum_{i=2}^{N_u} u_i \left(\frac{r}{r_c}\right)^i\right], & r \leq r_c, \\ 0, & r > r_c, \end{cases}
$$

with $N_u = 3$. The term $(1 - r/r_c)^2$ ensures that the pseudopotential goes smoothly to zero at $r = r_c$, and the component $u_1(1 + 2r/r_c)$ constrains the pseudopotential to have zero derivative at the origin. This ensures that the pseudo-wavefunction is smooth, easing the application of numerical methods.

To determine the coefficients $\{u_i\}$ we numerically solve the scattering problem, extract the scattering phase shift $\delta^{\mathrm{UTP}}(k)$, and then minimize the total squared error in the phase shift over all scattering wavevectors k,

$$
\langle\left|\delta^{\mathrm{UTP}}(k) - \delta^{\mathrm{cont}}(k)\right|^2\rangle =
$$
$$
\int_0^{k_F} \left|\delta^{\mathrm{UTP}}(k) - \delta^{\mathrm{cont}}(k)\right|^2 g(k/k_F)\mathrm{d}k,
$$

where the weighting is given by the density of states in the center of mass frame $g(k) = k(4 - \frac{8}{\pi}[k\sqrt{1-k^2} + \arcsin(k)])$ [32]. An example UTP is shown in Fig. 2, confirming that this construction gives smooth potentials. The scattering phase shift error from the UTP is shown in Fig. 3, demonstrating that the UTP construction creates pseudopotentials that are significantly more accurate than the Troullier–Martins pseudopotentials and soft disk potential, and some 200 times more accurate than the hard disk. This is achieved by the phase shift error from the UTP being optimized to be zero at three different wavevectors, as opposed to the single wavevector for the TM pseudopotentials.

B. Pseudopotentials for bound states

Pseudopotentials may also be constructed for particles in a bound state, with $E < 0$. In order to accurately imitate the contact potential, the pseudopotentials must reproduce the bound state energy of the contact potential $E_b = -(4/a^2)\exp(-2\gamma)$, and also must accommodate only one bound state. We first discuss the square well pseudopotential, which has been used in previous ultracold atomic gas calculations, and then again develop smooth pseudopotentials using the TM formalism. For bound states there is no quantity like the scattering phase shift that can be used to directly determine the quality of the pseudopotentials. We therefore demonstrate their accuracy in a two-body inhomogeneous system in Section IV.

1. Square well potential

The square well potential has the form

$$
V^{\mathrm{SW}}(r) = \begin{cases} -U, & r \leq R, \\ 0, & r > R. \end{cases}
$$

This potential may be made arbitrarily close to the bound state contact interaction by taking the well radius $R \to 0$ and depth $U \to \infty$. Decreasing R, however, reduces the sampling efficiency and thereby increases the computational cost. We require R to be less than the average interparticle separation $\sim 1/k_F$, in order to avoid the unphysical situation of three or more particles interacting simultaneously. In Section IV we investigate the R dependence of the accuracy of the square well pseudopotential.

Because there is no analogue of the scattering phase shift for the bound system it is not possible to uniquely define a U and R for a given a, as we did for the soft disk potential in which U and R were related by the second order term in the expansion. However one parameter may be determined by ensuring that the bound state energy of the potential is E_b, and for a given R the value of U this sets can be found as a solution to

$$
\frac{-J_0(k_1 R)}{k_1 J_1(k_1 R)} = \frac{J_0(ik_2 R) - iY_0(-ik_2 R)}{k_2\left(-iJ_1(ik_2 R) + Y_1(-ik_2 R)\right)},
$$

where $k_1 = \sqrt{U - |E_b|}$ and $k_2 = \sqrt{|E_b|}$. An example of a square well potential is shown in Fig. 4. Except in the limit of being infinitely deep and narrow, the square well potential does not give rise to the same wavefunction as the true contact interaction, but within the potential the wavefunction and therefore probability density is too small. This means that in the presence of an external potential (for example an harmonic trap, as in Section IV) there is too much weight at large particle separations, giving rise to inaccurate values of the system's energy. As $R \to 0$ the wavefunction approaches the exact form given by Equation (5).

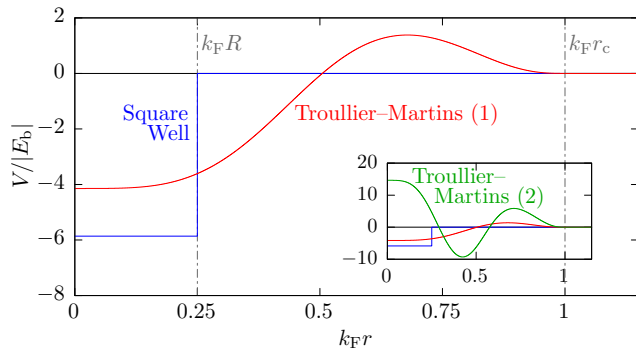FIG. 4. (Color online) Bound state pseudopotentials for the contact potential with $k_{\mathrm{F}}a = 1/2$. A square well with $k_{\mathrm{F}}R = 1/4$ is shown in blue, and the TM (1) pseudopotentials with $k_{\mathrm{F}}r_{\mathrm{c}} = 1$ is shown in red. Inset: the TM (2) pseudopotential, shown in green and with $k_{\mathrm{F}}r_{\mathrm{c}} = 1$, behaves qualitatively differently near particle coalescence.

FIG. 5. Analytic energy levels for two particles in an harmonic trap as a function of the dimensionless interaction strength $g = -1/\ln(a/d)$. The excited states (solid lines) correspond to the scattering states of the contact potential and the ground state (dashed line) corresponds to the bound state. The non-interacting energies are shown by circles along the line $g = 0$.

2. Troullier–Martins pseudopotential

The Troullier–Martins pseudopotential resolves the problem of having too much weight at large particle separations by being a norm-conserving pseudopotential, and so has the correct amount of weight within and outside of its cutoff radius. The construction of the TM pseudopotential for the bound state is identical to that of the scattering state, except that the calibration energy is now given by the bound state energy $E_{\mathrm{c}} = E_{\mathrm{b}}$. The cutoff radius r_{c} should be kept smaller than the average interparticle separation $\sim 1/k_{\mathrm{F}}$ to reduce the probability of three or more particles interacting at once, but there is no lower bound on r_{c}: similarly to the case of the square well, reducing r_{c} increases the accuracy but also the computational cost of simulations. The square well and TM pseudopotentials are shown in Fig. 4.

One of the TM pseudopotentials, labeled (1) in Fig. 4, behaves as would be expected qualitatively for a short-ranged potential giving rise to a bound state: it has a large negative region near particle coalescence. The other TM solution, labeled (2) and shown in the inset to Fig. 4, does not show this behavior, instead having an attractive region at finite particle separation. This will give rise to a non-zero expected separation between bound particles, which is physically discordant with the contact interaction. We therefore reject the TM (2) pseudopotential because of its unphysical behavior and select the TM (1) pseudopotential instead, referring to it henceforth simply as the TM pseudopotential.

Since all particles in bound states have approximately the same energy, the UTP formalism does not offer any advantage in this system. We now move on to testing the pseudopotentials in an inhomogeneous two-body system.
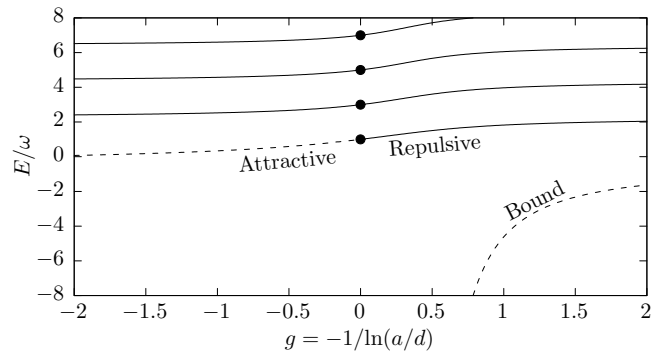
IV. TWO FERMIONS IN AN HARMONIC TRAP

We have constructed pseudopotentials that describe the scattering behavior of two isolated fermions. To test the pseudopotentials we turn to the experimentally realizable [38, 39] system of two distinguishable fermions in a circular harmonic trapping potential $\frac{1}{4}\omega^2 r^2$ of frequency ω. This system also has the advantage of being analytically soluble, which provides a stringent test for the pseudopotentials that we will use in many-body simulations.

A. Analytic energy levels

In the center-of-mass frame the Schrödinger Equation for two distinguishable fermions in an harmonic trap is given by

$$-\nabla^2\psi(r) + \frac{1}{4}\omega^2 r^2\psi(r) + V^{\mathrm{cont}}(r)\psi(r) = E\psi(r)$$

where the interparticle interaction term $V^{\mathrm{cont}}(r)\psi$ can be replaced by a boundary condition given by Equation (6). For the contact interaction the energy levels in the center-of-mass frame are solutions to the nonlinear equation [33, 40]

$$\Psi\left(-\frac{E}{2\omega} + \frac{1}{2}\right) = \ln\left(\frac{d^2}{a^2}\mathrm{e}^{-2\gamma}\right),\tag{10}$$

where $d = \sqrt{2/\omega}$ is the characteristic length scale of the trap and Ψ is the digamma function. These solutions are shown in Fig. 5 as a function of the dimensionless interaction strength $g = -1/\ln(a/d)$. In the non-interacting case $g = a = 0$ the energies have the expected values

$E = \omega(2n + 1)$ for non-interacting particles. As the repulsive interaction strength $g > 0$ in Fig. 5 increases, the energy increases and at $g \to \infty$ joins onto the energy of the attractive branch at $g \to -\infty$, in an analogue of unitarity in the BEC-BCS crossover [41]. The bound state of the contact potential survives in this inhomogeneous system as the deep bound state at $g > 0$.

B. Accuracy of the pseudopotentials

We compare the estimates of the center-of-mass energies of two particles in an harmonic trap to the analytic result in: Fig. 6(a), for repulsive interactions; Fig. 6(b) for attractive interactions; and Fig. 6(c) for bound particles.

In the repulsive case, we find that the hard and soft disk potentials and TM pseudopotential are accurate at small interaction strengths, but at large interaction strengths the error in the calculated energies is greater than 10%. The UTP pseudopotential is around 10 times more accurate at high interaction strengths, and becomes exact in the non-interacting limit.

To choose the radii of the potentials for the attractive and bound branches we follow the approach used in Ref. [30] and use a TM pseudopotential with a cutoff radius of $nr_c^2 = 10^{-2}$, where $n = \omega/2\pi$ is the peak density of two non-interacting particles in the trap. We compare this to square wells with radii given by the same $nR^2 = 10^{-2}$ and the smaller $nR^2 = 10^{-4}$ [29, 30]. We note that in both the attractive and bound branches, reducing the well radius increases the accuracy of the square well potential, but that the TM pseudopotential gives up to 10 times higher accuracy than a square well with a radius $1/10$ the size. The ability to use a larger cutoff radius with the TM pseudopotential brings significant benefits in numerical sampling of the potential, with the sampling efficiency expected to scale as $\sim r_c^2$. The increased accuracy can be related to the fact that the square well gives rise to wavefunctions with too much weight at large particle separations, raising the energy in the external trap, whilst the TM pseudopotential is norm-conserving, having the correct weight in the wavefunction outside r_c. The norm-conservation condition ensures that the TM pseudopotential gives a bound state wavefunction that is robust against changes in the local environment, and hence performs well in the spatially varying harmonic trap. As opposed to the single calibration energy of the TM pseudopotential, in constructing the UTP we would average over a range of energies. This would offer no advantage in the attractive and bound branches, where there is a definite binding energy for the pair of particles, and so we do not examine the UTP in these branches.

We have shown that for particles in an harmonic trap with attractive interactions, the TM pseudopotential gives an increase in both accuracy and sampling efficiency relative to the square well potential. For two
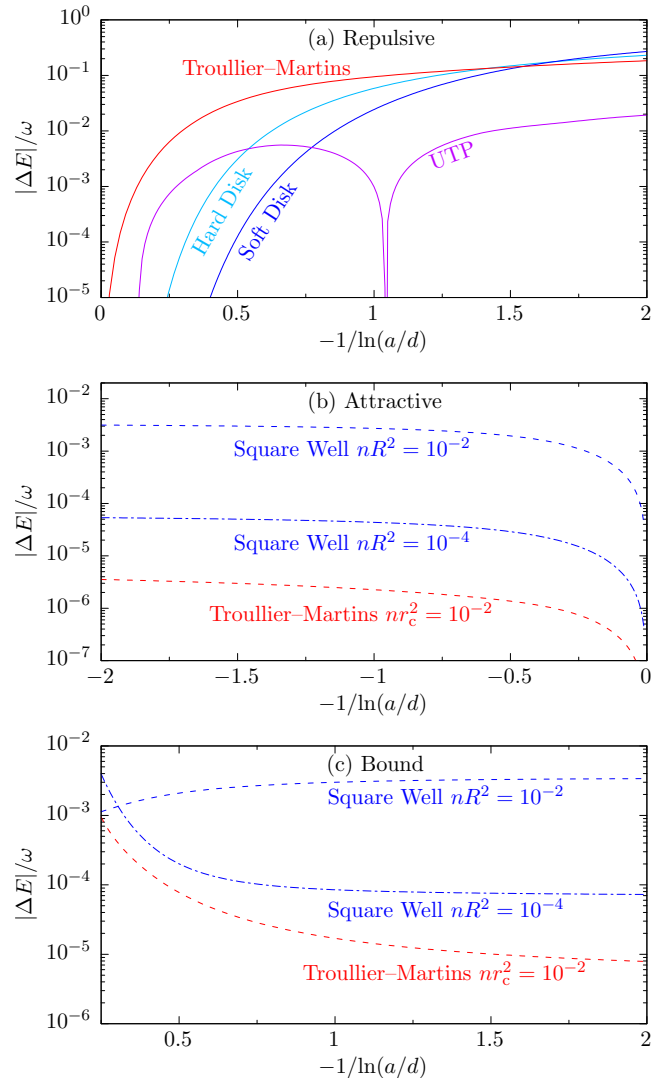

FIG. 6. (Color online) Error in the center-of-mass energy of two fermions in an harmonic trap calculated using pseudopotentials from the analytic value of the energy from Equation (10). (a) Error in the center-of-mass energy of particles with a repulsive interaction as a function of interaction strength for the hard disk in cyan, the soft disk in blue, the TM pseudopotential in red, and the UTP in magenta. (b) Error in the center-of-mass energy of particles with a weakly attractive interaction as a function of interaction strength. The square well pseudopotentials have radii given by $nR^2 = 10^{-4}$ and $nR^2 = 10^{-2}$, and the TM pseudopotential has a cutoff radius given by $nr_c^2 = 10^{-2}$, with the different cutoff radii denoted by different types of line dashing. (c) Error in the bound state energy as a function of interaction strength.

particles with repulsive interactions, the use of a UTP can offer a 10-fold increase in accuracy relative to using the TM pseudopotential or hard or soft disk potentials. We now go on to demonstrate the scaling benefits of the UTP in a many-body simulation.

V. FERMI GAS

Having demonstrated the effectiveness of the UTP for studying the two body scattering problem and two distinguishable fermions in an harmonic trap, we now demonstrate the advantages of the UTP in a prototypical setting: a two-dimensional homogeneous Fermi gas. Such a system serves as a benchmark for cold atom experiments [42, 43] and also as a model for electrons in conductors.

We focus on the repulsive branch of the contact interaction. Here the hard and soft disk potentials are uniquely defined for a given interaction strength, and may not be improved to attain arbitrarily high accuracy, as is possible in the attractive and bound branches by reducing the well radius to zero. This allows us to demonstrate the intrinsic benefits of the UTP formalism over the hard and soft disk potentials.

The smoothness of the UTP relative to the hard and soft disk potentials will be reflected in the many-body wavefunction, which will make it easy to work with numerically. Having shown in Section IV that the UTP is more accurate than the competing hard and soft disk potentials and Troullier–Martins pseudopotential, we proceed here to verify the accuracy of the UTP by comparing the energy of a Fermi gas with first- and second-order perturbation theory calculations [44–46].

A. Formalism

To calculate the ground state energies we use the diffusion Monte Carlo (DMC) technique. DMC is a highly-accurate Green's function projector method for determining ground state energies and expectation values [47–49], and it is well-suited to investigating homogeneous gaseous phases. We use the CASINO implementation [50] of the DMC method with a Slater–Jastrow trial wavefunction $\Psi = \mathrm{e}^J D_\uparrow D_\downarrow$, where $D_\uparrow$ $(D_\downarrow)$ is a Slater determinant of plane-wave states for the spin up (down) channel. The Jastrow factor e^J describes correlations between particles, with

$$J = \sum_{\substack{j \neq i \\ \alpha, \beta \in \{\uparrow, \downarrow\}}} \left(1 - \frac{r_{ij}}{L_\mathrm{c}}\right)^3 u_{\alpha\beta}\left(r_{ij}\right) \Theta\left(L_\mathrm{c} - r_{ij}\right), \quad (11)$$

where $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$ is the distance between two particles with labels i and j, and $u_{\alpha\beta}$ are eighth-order polynomials, whose parameters are optimized using variational Monte Carlo subject to the symmetry requirements $u_{\uparrow\uparrow} = u_{\downarrow\downarrow}$ and $u_{\uparrow\downarrow} = u_{\downarrow\uparrow}$. L_c is a cutoff length that we set equal to the radius of a circle inscribed within the simulation cell, and Θ is the Heaviside step function.

We calculate the ground state energy expectation value for 49 spin-up and 49 spin-down particles in a homogenous two-dimensional system for increasing interaction strengths $-1/\ln(k_\mathrm{F}a)$ up to a maximum value of 1.8 before the system would phase separate into a fully polar-

ized state. To accurately capture the hard disk wavefunction at small inter-particle distances in our DMC simulations we add an additional term to the Jastrow factor in Equation (11),

$$u_\mathrm{H}(r) = \begin{cases} -\infty, & r \leq R, \\ \log[\tanh(\frac{r/R - 1}{1 - r/L_\mathrm{c}})], & R < r < L_\mathrm{c}, \\ 0, & r \geq L_\mathrm{c}, \end{cases} \quad (12)$$

as in Ref. [11], where R is the hard disk radius. In the present study the additional term applies to opposite spins only.

We extrapolate to zero DMC timestep to obtain accurate ground state energies. For each data point we run three simulations with timesteps $0.25dt, 0.5dt, dt$, [51] with dt the maximum timestep in the linear regime, and extrapolate to zero timestep by minimizing the weighted least squares fit. All error bars represent the DMC stochastic error combined with the concomitant uncertainty in the timestep extrapolation. We expect that the use of a quadratic DMC algorithm would give similar results [52, 53].

B. Results

In Fig. 7 we compare ground state energies of the Fermi gas obtained using the different potentials. It is clear that for $-1/\ln(k_\mathrm{F}a) > 0.7$ both the hard and soft disk potentials, as well as the Troullier–Martins pseudopotential, are insufficient to obtain the desired $10^{-4}E_\mathrm{F}$ accuracy that has been obtained in other DMC studies of homogeneous systems [30–32, 54].

To verify the DMC results we compare our estimates for the ground state energy with perturbation theory [44–46]. As can be seen in Fig. 7(b), first order perturbation theory $E^{(1)} = \frac{E_\mathrm{F}}{2}(1 + [-1/\ln(k_\mathrm{F}a)])$ deviates quadratically in the interaction strength $-1/\ln(k_\mathrm{F}a)$ from the UTP result as expected, and second order perturbation theory

$$E^{(2)} = \frac{E_\mathrm{F}}{2}\left[1 + \left(\frac{-1}{\ln(k_\mathrm{F}a)}\right) + \left(\frac{3}{4} - \ln(4\mathrm{e}^{-\gamma})\right)\left(\frac{-1}{\ln(k_\mathrm{F}a)}\right)^2\right]$$

deviates cubically in $-1/\ln(k_\mathrm{F}a)$ and outperforms first order perturbation theory. In Fig. 7 we also show the result obtained in Ref. [46] using a partial resummation of Feynman diagrams in the Galitskii-Feynman (GF) scheme which is correct to order $\mathcal{O}\left([-1/\ln(k_\mathrm{F}a)]^3\right)$, and note that this indeed deviates cubically in interaction strength from the UTP result. The agreement of the scaling behavior of the energy calculated using the UTP with interaction strength when compared to these analytic results confirms the accuracy of the UTP.

In addition to the analytic approximations, we compare our DMC results with an independent study using the hard disk potential and the same number of particles in Ref. [22], labeled GB. We note that their predicted energies are higher than those from our DMC calculations

FIG. 8. (Color online) Standard deviation of the local energy distribution of the trial wavefunction. The hard and soft disk pseudopotentials exhibit a larger standard deviation, due to the sudden changes in energy when two particles approach one another.
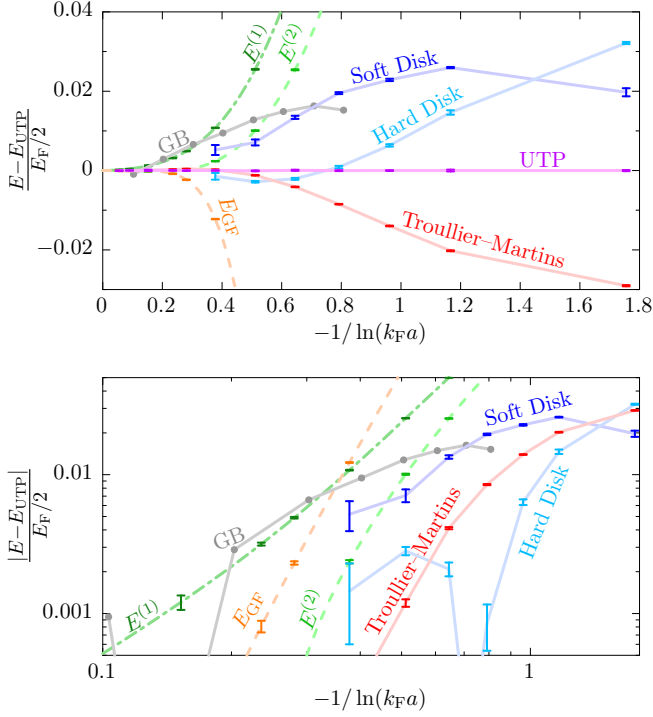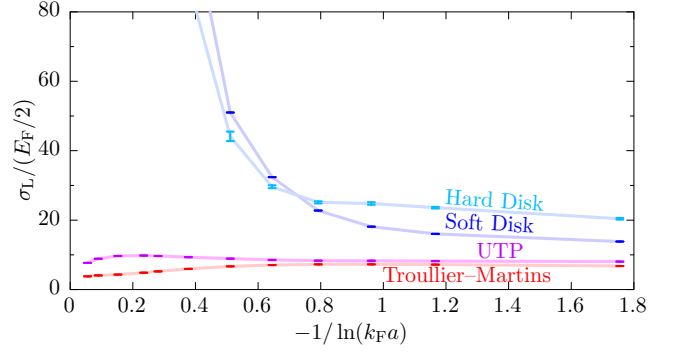
FIG. 7. (Color online) (Top) Differences in ground state energy from the result obtained with the UTP as a function of interaction strength, normalized by the energy of the non-interacting system. The green lines denoted $E^{(1)}$ and $E^{(2)}$ are predictions from first- and second-order perturbation theory [44–46] and E_{GF} is the result of a Galitskii-Feynman partial resummation of Feynman diagrams reported in Ref. [46], shown in orange. GB is the Monte Carlo result from Ref. [22], calculated using a hard disk potential and shown in gray, and our results using the hard disk are shown in cyan, the soft disk in blue, the TM pseudopotential in red, and the UTP in magenta. (Bottom) The same results on a logarithmic scale.

using the hard disk potential, and as DMC is a variational method this indicates that our trial wavefunction is likely more accurate than was available to the authors of Ref. [22], possibly due to our inclusion of a Jastrow factor with variational parameters.

Having confirmed the accuracy of the UTP we now examine its performance benefits. The local energy, $E_L = \Psi^{-1}\hat{H}\Psi$, is a crucial quantity in DMC calculations [50]. The stochastic error in a DMC calculation is proportional to the standard deviation σ_L in the local energy distribution, and therefore a smoother local energy will give rise to more accurate results for the same computation time. Fig. 8 shows the standard deviation of the local energy distribution of the trial wave function when using all of our pseudopotentials. Both the UTP and TM pseudopotentials benefit from their smoothness in obtaining a lower local energy standard deviation compared to the hard and soft disk potentials. For weak interactions the hard disk potential benefits from an additional Jastrow factor term, Equation (12), relative to the soft disk po-

tential, whose height U also diverges as $a \to 0$. However for larger interactions the soft disk potential results in a smoother wavefunction than the hard disk potential and therefore has lower local energy variance. The variance in the local energies diverges for the hard and soft disk potentials for weak interactions, whereas it decays for the UTP and TM pseudopotentials. The standard deviation for the TM pseudopotential is slightly lower than the UTP at all interaction strengths, which is understood from the larger size of the potential for the UTP in Fig. 2 compared to the TM pseudopotential. This behavior is similar to the 3D case reported in Ref. [30].

The reduced variance in the local energy lowers the computational effort T required for a DMC calculation, which scales as $T \propto \sigma_L^2/dt$ [31, 32, 55]. From Fig. 8 we see that at intermediate interaction strength $-1/\ln(k_F a) = 0.8$ the variance of the local energy for the UTP is 2.7 and 3.0 times lower than for the soft and hard disk potentials respectively, corresponding to a speedup of 7.5 and 9.1.

In addition to the lower local energy variance, our pseudopotentials offer an additional speedup. The DMC estimate of the energy must be extrapolated to zero timestep, and the larger the region of linear dependence of energy on timestep, the larger timestep can be used. This reduces computational effort even further, as $T \propto 1/dt$. In Fig. 9 we observe that the extent of the linear regime of the error in ground state energy with timestep differs between the pseudopotentials: it extends up to $dt_{HD} = 1.25 \times 10^{-3}/E_F$ for the hard disk, up to $dt_{SD} = 2.5 \times 10^{-3}/E_F$ for the soft disk, up to $dt_{UTP} = 5.0 \times 10^{-3}/E_F$ for the UTP, and up to $dt_{TM} = 1.0 \times 10^{-2}/E_F$ for the TM pseudopotential. This means that the maximum timestep for a calculation with the UTP is two and four times larger than for the soft and hard disk potentials respectively. Combining this with the reduced variance we therefore accomplish a total speedup of at least 15 times by using the UTP instead of the hard and soft disks.

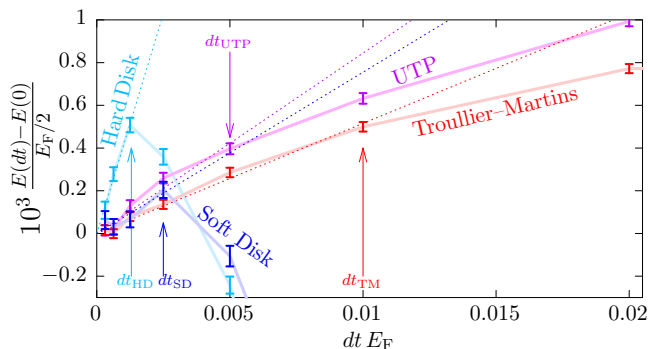To summarize, we have demonstrated the importance

FIG. 9. (Color online) Error in the estimated ground state energy as a function of DMC timestep. The results using the hard disk are shown in cyan, the soft disk in blue, the TM potential in red and the UTP in magenta, with solid lines indicating the values calculated using DMC and dotted lines a linear extrapolation. This enables the identification of when the error leaves the linear regime.

of using a pseudopotential with scattering properties that accurately describe the contact interaction. For weak interactions we observe that a divergence in the variance in the local energy severely constrains the accuracy of DMC simulations with soft or hard disk potentials. At strong interactions these inaccurate potentials introduce a significant bias into the results, such that we were unable to attain the $10^{-4}E_{\mathrm{F}}$ target accuracy in the ground state energy. However the UTP delivers highly accurate results over the full range of interaction strengths and additionally offers 15 times better computational performance. We therefore recommend the UTP as an accurate and efficacious tool for studying the contact interaction in 2D.

VI. DISCUSSION

We have developed a high-accuracy pseudopotential for the contact interaction in 2D, building on the work of Ref. [30]. We have demonstrated that our ultratransferable pseudopotential provides accurate scattering phase shifts, accurate energies for two harmonically confined particles, and we have demonstrated its advantages in many-body simulations. The energies obtained with our

UTP are over 10 times more accurate in the repulsive branch of the interaction than is afforded by the hard and soft disk potentials used in recent studies. Moreover, we have demonstrated that for many-body systems our pseudopotential delivers a speedup of at least 15 times in diffusion Monte Carlo computations, on top of the more accurate result.

The performance and ease of construction of the pseudopotential suggests that it could be widely applicable across first-principles methods beyond quantum Monte Carlo. The pseudopotential formalism has already been used to study the Coulomb [31] and dipolar [32] interactions. Although in this work we have focused on using the pseudopotential to accurately capture the scattering properties of the contact interaction, our formalism allows the further improvement of modeling of quantum gases by calibrating the pseudopotentials to more accurately describe the scattering properties of the underlying Feshbach resonance interaction. To next lowest order in scattering wavevector, this corresponds to including the effective range term essential for describing narrow Feshbach resonances, which may exhibit exotic breached superfluidity [56, 57], or other interactions with non-zero effective ranges, which are applicable in the study of nucleon reactions [58]. Rather than a description in terms of the scattering phase shift, the pseudopotentials could instead be calibrated to other scattering properties. For example, they could be calibrated to the cross-section for elastic scattering measured experimentally via the thermalization rate, or the inelastic loss coefficient, to capture the full physical interaction between particles [20].

[1] Y.J. Lin, R.L. Compton, K. Jiménez-García1, J.V. Porto, and I.B. Spielman, Nature **462**, 628 (2009).
[2] P.A. Lee, N. Nagaosa, and X.G. Wen, Reviews of Modern Physics **78**, 17 (2006).
[3] S. Sachdev, Nature Physics **4**, 173 (2008).
[4] M.Z. Hasan and C.L. Kane, Reviews of Modern Physics **82**, 3045 (2010).
[5] G.J. Conduit, Phys. Rev. A **82**, 043604 (2010).
[6] V. Ngampruetikorn, J. Levinsen, and M.M. Parish, Phys. Rev. Lett. **111**, 265301 (2013).
[7] K. Martiyanov, V. Makhalov, and A. Turlapov, Phys. Rev. Lett. **105**, 030404 (2010).
[8] A.T. Sommer, L.W. Cheuk, M.J.H. Ku, W.S. Bakr, and M.W. Zwierlein, Phys. Rev. Lett. **108**, 045302 (2012).
[9] K. Günter, T. Stöferle, H. Moritz, M. Köhl, and T. Esslinger, Phys. Rev. Lett. **95**, 230401 (2005).
[10] H. Shi, S. Chiesa, and S. Zhang, Phys. Rev. A **92**, 033603 (2015).

[11] N.D. Drummond, N.R. Cooper, R.J. Needs, and G.V. Shlyapnikov, Phys. Rev. B **83**, 195429 (2011).

[12] J.A. Bert, B.Kalisky, C. Bell, M. Kim, Y. Hikita, H.Y. Hwang, and K.A. Moler, Nature Physics **7**, 767 (2011).

[13] A. Görlitz, J.M. Vogels, A.E. Leanhardt, C. Raman, T.L. Gustavson, J.R. Abo-Shaeer, A.P. Chikkatur, S. Gupta, S. Inouye, T. Rosenband, and W. Ketterle, Phys. Rev. Lett. **87**, 130402 (2001).

[14] K. Fenech, P. Dyke, T. Peppler, M.G. Lingham, S. Hoinka, H. Hu, and C.J. Vale, Phys. Rev. Lett. **116**, 045302 (2016).

[15] I. Boettcher, L. Bayha, D. Kedar, P.A. Murthy, M. Neidig, M.G. Ries, A.N. Wenz, G. Zürn, S. Jochim, and T. Enss, Phys. Rev. Lett. **116**, 045303 (2016).

[16] R.A. Hart, P.D. Duarte, T.-L. Yang, X. Liu, T. Paiva, E. Khatami, R.T. Scalettar, N. Trivedi, D.A. Huse, and R.G. Hulet, Nature **519**, 211 (2015).

[17] M.H. Anderson, J.R. Ensher, M.R. Matthews, C.E. Wieman, and E.A. Cornell, Science **269**, 198 (1995).

[18] C.A. Regal, M. Greiner, and D.S. Jin, Phys. Rev. Lett **92**, 040403 (2004).

[19] M.W. Zwierlein, C.A. Stan, C.H. Schunck, S.M.F. Raupach, A.J. Kerman, and W. Ketterle, Phys. Rev. Lett **92**, 120403 (2004).

[20] C. Chin, R. Grimm, P. Julienne and E. Tiesinga, Reviews of Modern Physics **82(2)**, 1225 (2010).

[21] L. Xing, Phys. Rev. B **42**, 8426 (1990).

[22] G. Bertaina, The European Physical Journal Special Topics **217**, 153 (2013).

[23] M. Casula, D.M. Ceperley, and E.J. Mueller, Phys. Rev. A **78**, 033607 (2008).

[24] X. Li, J. Kolorenč, and L. Mitas, Phys. Rev. A **84**, 023615 (2011).

[25] G. Bertaina, E. Fratini, S. Giorgini, and P. Pieri, Phys. Rev. Lett. **110**, 115303 (2013).

[26] S.Q. Zhou, D.M. Ceperley, and S. Zhang, Phys. Rev. A **84**, 013625 (2011).

[27] S. Nascimbène, N. Navon, K.J. Jiang, F. Chevy, and C. Salomon, Nature **463**, 1057 (2010).

[28] G. Bertaina and S. Giorgini, Phys. Rev. Lett. **106**, 110403 (2011).

[29] G.E. Astrakharchik, J. Boronat, J. Casulleras, and S. Giorgini, Phys. Rev. Lett. **93**, 200404 (2004).

[30] P.O. Bugnion, P. López Ríos, R.J. Needs, and G.J. Conduit, Phys. Rev. A **90**, 033626 (2014).

[31] J.H. Lloyd-Williams, R.J. Needs, and G.J. Conduit, Phys. Rev. B **92**, 075106 (2015).

[32] T.M. Whitehead and G.J. Conduit, arXiv:1601.07746.

[33] X.J. Liu, H. Hu, and P.D. Drummond, Phys. Rev. B **82**, 054524 (2010).

[34] A. Farrell and P.B. van Zyl, Can. J. Phys. **88**, 817 (2010).

[35] N. Troullier and J.L. Martins, Phys. Rev. B **43**, 1993 (1991).

[36] We provide a Mathematica notebook to generate the hard and soft disk potentials, square well potential, Troullier–Martins pseudopotentials for both attractive and bound states, and UTP at `www.repository.cam.ac.uk/handle/1810/253582`.

[37] G.J. Conduit, Phys. Rev. B **87**, 184414 (2013).

[38] S. Murmann, A. Bergschneider, V.M. Klinkhamer, G. Zürn, T. Lompe, S. Jochim, Phys. Rev. Lett. **114**, 080402 (2015).

[39] G. Zürn, F. Serwane, T. Lompe, A.N. Wenz, M.G. Ries, J.E. Bohn, and S. Jochim, Phys. Rev. Lett. **108**, 075303 (2012).

[40] T. Busch, B.G. Englert, K. Rzazewski and M. Wilkens, Foundations of Physics **28**, 549 (1998).

[41] M.G. Ries, A.N. Wenz, G. Zürn, L. Bayha, I. Boettcher, D. Kedar, P.A. Murthy, M. Neidig, T. Lompe, and S. Jochim, Phys. Rev. Lett **114**, 230401 (2015).

[42] B. Fröhlich, M. Feld, E. Vogt, M. Koschorreck, W. Zwerger, and M. Köhl, Phys. Rev. Lett. **106**, 105301 (2011).

[43] K. Martiyanov, V. Makhalov, and A. Turlapov, Phys. Rev. Lett. **105**, 030404 (2010).

[44] J.R. Engelbrecht and M. Randeria, Phys. Rev. B **45**, 12419 (1992).

[45] J.R. Engelbrecht, M. Randeria, and L. Zhang, Phys. Rev. B **45**, 10135 (1992).

[46] L. He, Phys. Rev. A **90**, 053633 (2014).

[47] D.M. Ceperley, and B.J. Alder, Phys. Rev. Lett. **45**, 566 (1980).

[48] W.M.C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Reviews of Modern Physics **73**, 33 (2001).

[49] C. Umrigar, The Journal of Chemical Physics **99**, 2865 (1993).

[50] R.J. Needs, M.D. Towler, N.D. Drummond, and P. López Ríos, Journal of Physics: Condensed Matter **22**, 023201 (2010).

[51] R.M. Lee, G.J. Conduit, N. Nemec, P. López Ríos, and N.D. Drummond, Phys. Rev. E **83**, 066706 (2011).

[52] M. Mella, G. Morosi, and D. Bressanini, Phys. Rev. E **61**, 2050 (2000).

[53] A. Sarsa, J. Boronat, and J. Casulleras, J. Chem. Phys. **116**, 5956 (2002).

[54] G.J. Conduit, A.G. Green, and B.D. Simons, Phys. Rev. Lett. **103**, 207201 (2009).

[55] A. Ma, N.D. Drummond, M.D. Towler, and R.J. Needs, Phys. Rev. E **71**, 066704 (2005).

[56] W.V. Liu and F. Wilczek, Phys. Rev. Lett. **90**, 047002 (2003).

[57] M.M. Forbes, E. Gubankova, W.V. Liu, and F. Wilczek, Phys. Rev. Lett **94**, 017001 (2005).

[58] V.A. Babenko and N.M. Petrov, Physics of Atomic Nuclei **76**, 684 (2013).

[59] T.M. Whitehead, L.M. Schonenberg, N. Kongsuwan, R.J. Needs, and G.J. Conduit, Cambridge University DSpace Repository, `www.repository.cam.ac.uk/handle/1810/254305`.