

AN INVERSE ITERATION METHOD FOR EIGENVALUE PROBLEMS WITH EIGENVECTOR NONLINEARITIES*

ELIAS JARLEBRING[†], SIMEN KVAAL[‡], AND WIM MICHIELS[§]

Abstract. Consider a symmetric matrix $A(v) \in \mathbb{R}^{n \times n}$ depending on a vector $v \in \mathbb{R}^n$ and satisfying the property $A(\alpha v) = A(v)$ for any $\alpha \in \mathbb{R} \setminus \{0\}$. We will here study the problem of finding $(\lambda, v) \in \mathbb{R} \times \mathbb{R}^n \setminus \{0\}$ such that (λ, v) is an eigenpair of the matrix $A(v)$ and we propose a generalization of inverse iteration for eigenvalue problems with this type of eigenvector nonlinearity. The convergence of the proposed method is studied and several convergence properties are shown to be analogous to inverse iteration for standard eigenvalue problems, including local convergence properties. The algorithm is also shown to be equivalent to a particular discretization of an associated ordinary differential equation, if the shift is chosen in a particular way. The algorithm is adapted to a variant of the Schrödinger equation known as the Gross–Pitaevskii equation. We use numerical simulations to illustrate the convergence properties, as well as the efficiency of the algorithm and the adaptation.

Key words. nonlinear eigenvalue problems, inverse iteration, Gross–Pitaevskii equation, convergence factors

AMS subject classifications. 65F15, 65H17, 35P30, 35Q55, 65Z05

DOI. 10.1137/130910014

1. Introduction. Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ be a symmetric matrix depending on a vector. We will consider the corresponding nonlinear eigenvalue problem where the vector-valued parameter of A , denoted $v \in \mathbb{R}^n$, equals an eigenvector of the symmetric matrix $A(v)$. That is, we wish to find $(\lambda, v) \in \mathbb{R} \times \mathbb{R}^n \setminus \{0\}$ such that

$$(1.1) \quad A(v)v = \lambda v,$$

and we will call λ an eigenvalue, v an eigenvector, and (λ, v) an eigenpair of (1.1). In this work we will assume that A is three times continuously differentiable with respect to v for any $v \neq 0$. Note that we have assumed that $A(v)$ is symmetric in order to work in real formulation. By separating real and imaginary parts, we will, however, in the examples, show that we can approach hermitian complex problems.

Moreover, we shall assume that A satisfies

$$(1.2) \quad A(\alpha v) = A(v) \quad \text{for any } \alpha \in \mathbb{R} \setminus \{0\},$$

such that the solution is independent of the scaling of v , i.e., if (λ, v) is a solution to (1.1), then $(\lambda, \alpha v)$ is also a solution for any $\alpha \in \mathbb{R} \setminus \{0\}$. If a problem does not have property (1.2), but instead satisfies a normalization constraint, we can often transform it to an equivalent problem with property (1.2). For instance, consider

*Submitted to the journal's Methods and Algorithms for Scientific Computing section February 19, 2013; accepted for publication (in revised form) June 17, 2014; published electronically August 19, 2014.

<http://www.siam.org/journals/sisc/36-4/91001.html>

[†]Dept. of Mathematics, NA group, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden (eliasj@kth.se).

[‡]Centre for Theoretical and Computational Chemistry, Department of Chemistry, University of Oslo, P.O. Box 1033 Blindern, N-0315 Oslo, Norway (simen.kvaal@kjemi.uio.no).

[§]K.U. Leuven, Department of Computer Science, 3001 Heverlee, Belgium (wim.michiels@cs.kuleuven.be).

TABLE 1
Applicability of the convergence characterizations in this paper.

	Shift σ	Error $\ v_k - v_*\ $	Characterization
Section 3	arbitrary	$\ v_k - v_*\ \ll 1$	One step of (1.3) satisfies the first-order expansion (1.5)
Section 4	$\sigma \ll \lambda_*$	arbitrary	One step of (1.3) approximates the trajectory of the ODE (1.6)

$\tilde{A}(\cdot)$ which is symmetric with respect to v , i.e., $\tilde{A}(v) = \tilde{A}(-v)$. Also suppose \tilde{A} does not satisfy (1.2). The problem to find a *normalized* vector $x \in \mathbb{R}^n$, i.e., $\|x\| = 1$, such that $\tilde{A}(x)x = \lambda x$ is equivalent to (1.1) if we define $A(v) := \tilde{A}(v/\|v\|)$ such that A does satisfy (1.2).

In this paper we propose a new algorithm for (1.1), which is a generalization of inverse iteration for standard linear eigenvalue problems. More precisely, we consider the iteration,

$$(1.3) \quad v_{k+1} = \alpha_k (J(v_k) - \sigma I)^{-1} v_k,$$

where $\alpha_k = 1/\|(J(v_k) - \sigma I)^{-1} v_k\|$ and J is the Jacobian of the left-hand side of (1.1) with respect to v , i.e.,

$$(1.4) \quad J(v) := \frac{\partial}{\partial v} (A(v)v) \in \mathbb{R}^{n \times n}.$$

The scalar $\sigma \in \mathbb{R}$ is called the shift and can be used to control to which eigenvalue the iteration converges, and has great influence on the speed of convergence.

The iteration (1.3) can be interpreted as a generalization of inverse iteration, which is one of the most used algorithms for eigenvalue computations. Convergence analysis can be found in [27, 14] and in the studies of Rayleigh quotient iteration in the classical series of works of Ostrowski, e.g., [26]. There are also inexact versions of inverse iteration and its variant Rayleigh quotient iteration, e.g., [28, 9]. These results are often used in combination with preconditioning techniques for inverse iteration [24, 23, 17]. Inverse iteration has also been generalized to eigenvalue problems with eigenvalue nonlinearities, e.g., [22, 20, 29] for which convergence has been studied in [15, 16]. There is an algorithm for problems with eigenvector nonlinearities and based on solving linear systems in [21]. See also the results on eigenvector nonlinearities in [7]. To our knowledge, the iteration (1.3) for problems of the type (1.1) has not been presented or analyzed in the literature.

We characterize the convergence of the iteration (1.3) in two ways, which lead to conclusions for the behavior in two situations; when the error $\|v_k - v_*\|$ is small or when the shift satisfies $\sigma \ll \lambda_*$; see Table 1.

In particular, in the local convergence analysis (in section 3) we show the following. For any eigenpair (λ_*, v_*) of (1.1), the iteration satisfies

$$(1.5) \quad v_{k+1} \pm v_* = |\lambda_* - \sigma| F_*(v_k - v_*) + O(\|v_k - v_*\|^2),$$

where the sign depends on $\text{sign}(\lambda_* - \sigma)$, and we derive an expression for $F_* \in \mathbb{R}^{n \times n}$. If the shift is sufficiently close to the eigenvalue, the iteration is locally convergent. The convergence is in general linear and the convergence factor is proportional to the distance between the shift and the eigenvalue. We also briefly discuss updating strategies for the shift and connections with Newton's method.

We also provide a characterization of (1.3) by deriving a relation with an ODE applicable in the situation when $\sigma \ll \lambda_*$. In particular, in section 4 we show the following. One step of (1.3) is equivalent to a particular type of discretization of the ODE

$$(1.6) \quad y'(t) = p(y(t))y(t) - A(y(t))y(t),$$

where p is the Rayleigh quotient

$$(1.7) \quad p(y) := \frac{y^T A(y)y}{y^T y},$$

if the shift is chosen such that $\sigma \ll \lambda_*$. The stationary solutions of (1.6) are solutions to (1.1). A small step length in the discretization corresponds to $\sigma \ll \lambda_*$, implying that one step of (1.3) approximates the trajectory of (1.6) if σ is chosen sufficiently negative. In the linear case $A(v) = B$, when the matrix B has one simple dominant eigenvalue (eigenvalue closest to $-\infty$), the ODE only has one stable stationary point (which corresponds to the dominant eigenvalue) and it is convergent for any starting value. If a similar situation occurs for the nonlinear case, i.e., the ODE (1.6) is convergent and only has one stable stationary point, then the iteration (1.3) is expected to converge to this solution for sufficiently negative σ and the convergence is expected to be independent of the starting vector. In particular, if the problem is close to linear, the iteration is expected to converge to the dominant eigenvalue.

The idea we use in section 4, basing the reasoning on interpreting the iterative method as a discretization or realization of an ODE, has been used in a number of other settings for eigenvalue problems before, e.g., for the QR-method [30, 31, 6], preconditioning techniques [17], and characterizations of the Rayleigh quotient iteration [19]. In fact, the ODE (4.1) is a nonlinear variant of the Oja flow [32]. See also the collection [3] and the description of iterations on manifolds in [8].

The algorithm is illustrated by showing how it can be adapted to the Gross–Pitaevskii equation (GPE) in section 5. The GPE is a standard model for particles in the state of matter called the Bose–Einstein condensate. See [18] and references in [1] for literature on the GPE. We discretize the GPE and transform it to the form (1.1). We also show how the solution to the linear system $(J(v_k) - \sigma I)^{-1}v_k$ can be found efficiently using the Sherman–Morrison–Woodbury formula. It turns out that in this setting, the ODE (1.6) is directly related to the technique called *imaginary time integration* [1]. This connection allows us to use results known for imaginary time integration, in particular, that the ODE always converges to a stationary solution and the experience presented in results in the literature indicate that this is often a physically relevant solution, e.g., the ground state. We further study the ODE and derive a heuristic choice of the step length h or, equivalently, a heuristic choice for the shift σ , which follows the trajectory of the ODE to sufficient accuracy, and still maintains a fast asymptotic convergence rate.

Although we are not aware of algorithms for (1.1), a number of successful algorithms do exist for the specific application we have in mind, i.e., the GPE in section 5. Besides the aforementioned methods based on imaginary time integration, there are methods for the GPE based on minimization [2, 4] as well as a Newton–Raphson approach [5]. Our approach is different in character since the reasoning stems from an eigenvalue algorithm (inverse iteration), and it has a different generality setting. It can be interpreted as a discretization of the ODE, but with an integration scheme and step length which we believe has not been used for the imaginary time integration of

the GPE. For completeness we also present results for another generalization of the algorithm which turns out to be essentially equivalent to an approach presented in [1] and can be seen as (1.3) where J is replaced by A .

The notation in this paper is mostly standard. We use $\frac{\partial}{\partial v}q(v)$ to denote the Jacobian of a vector or scalar q , i.e., if $v \in \mathbb{R}^n$ and $q(v) \in \mathbb{R}^k$, then $\frac{\partial}{\partial v}q(v) = (\frac{\partial}{\partial v_1}q(v), \dots, \frac{\partial}{\partial v_n}q(v)) \in \mathbb{R}^{k \times n}$. We use the notation $(\cdot)_{x=y}$ to denote substitution of x with y for the formula inside the parenthesis. As usual, the expression $\|Z\|$ denotes the Euclidean norm if Z is a vector and the spectral norm if Z is a matrix. The set of eigenvalues of a matrix $B \in \mathbb{R}^{n \times n}$ will be denoted $\lambda(B)$ and a dominant eigenvalue will be used to refer to an eigenvalue $\mu \in \lambda(B)$ for which all other eigenvalues have larger (or equal) real part, i.e., if $\mu \in \lambda(B)$ is a dominant eigenvalue, then any $\mu_2 \in \lambda(B)$ satisfies $\text{Re}(\mu) \leq \text{Re}(\mu_2)$.

2. Preliminaries and fixed point formulation. Throughout this paper we will in several situations use the following consequences of the scaling invariance property (1.2).

LEMMA 2.1 (scaling invariance). *Consider a scaling invariant function $Q : \mathbb{R}^n \rightarrow \mathbb{R}^{k \times \ell}$, i.e., $Q(\alpha v) = Q(v)$ for any $\alpha \in \mathbb{R} \setminus \{0\}$. Then, for any vectors $u \in \mathbb{R}^\ell$, $v \in \mathbb{R}^n$,*

$$(2.1) \quad \left(\frac{\partial}{\partial v}(Q(v)u) \right) v = 0.$$

In particular, given $A : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ satisfying the scaling invariance (1.2) and J defined by (1.4), we have

$$(2.2) \quad J(u)u = A(u)u$$

for any $u \in \mathbb{R}^n$. Moreover, if $u = v_$, when (λ_*, v_*) is an eigenpair of (1.1), then λ_* is an eigenvalue of $J(v_*)$ with eigenvector v_* .*

Proof. The left-hand side of (2.1) can be interpreted as a directional derivative in the direction of v . We have

$$\left(\frac{\partial}{\partial v}(Q(v)u) \right) v := \lim_{\varepsilon \rightarrow 0} \frac{Q(v + \varepsilon v)u - Q(v)u}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \frac{Q((1 + \varepsilon)v)u - Q(v)u}{\varepsilon} = 0,$$

which shows (2.1). The identity (2.2) follows from the chain rule. \square

The iteration can naturally be represented in a fixed point form

$$v_{k+1} = \varphi(v_k),$$

where

$$(2.3) \quad \varphi(v) := \frac{1}{\|\psi(v)\|} \psi(v)$$

and

$$(2.4) \quad \psi(v) := (J(v) - \sigma I)^{-1}v.$$

Obviously, $v_{k+1} = \varphi(v_k)$ when v_k is given by (1.3). It also turns out that the fixed points of φ are solutions to (1.1). However, the converse is not true. We shall now show that if $\sigma > \lambda_*$, we have $\varphi(v_*) = -v_*$, implying that the iterates v_k alternate between v_* and $-v_*$ if $v_0 = v_*$. An equivalence with the solutions to (1.1) and the vectors that satisfy $\varphi(v_*) = \pm v_*$, for some choice of the sign, can be achieved if we take the alternation into account, as can be seen as follows.

PROPOSITION 2.2 (equivalence of fixed points and eigenvectors).

- (a) Suppose (λ_*, v_*) is a solution to (1.1) with $\|v_*\| = 1$ and suppose $\sigma \neq \lambda_*$. Then, the fixed point map φ satisfies

$$(2.5) \quad \varphi(v_*) = \pm v_*.$$

The sign in (2.5) should be chosen positive if $\sigma < \lambda_*$, otherwise negative.

- (b) Suppose $v_* = \varphi(v_*)$ or $v_* = -\varphi(v_*)$. Then, (λ_*, v_*) is a solution to (1.1) with

$$(2.6) \quad \lambda_* := \sigma \pm \frac{1}{\|(J(v_*) - \sigma I)^{-1}v_*\|} = v_*^T A(v_*)v_*.$$

Proof. In order to show (a), suppose (λ_*, v_*) is an eigenpair and note that from Lemma 2.1 we have that

$$\lambda_* v_* = A(v_*)v_* = J(v_*)v_*.$$

By subtracting σv_* from both sides and subsequently multiplying both sides from the left by $\frac{1}{\lambda_* - \sigma}(J(v_*) - \sigma I)^{-1}$ we can simplify $\psi(v_*) = (J(v_*) - \sigma I)^{-1}v_* = \frac{1}{\lambda_* - \sigma}v_*$, and we consequently find that

$$(2.7) \quad \frac{1}{\|\psi(v_*)\|} = \frac{1}{\sqrt{\psi(v_*)^T \psi(v_*)}} = |\lambda_* - \sigma|.$$

Hence, φ evaluated at $v = v_*$ is explicitly

$$\varphi(v_*) = \frac{1}{\sqrt{\psi(v_*)^T \psi(v_*)}} \psi(v_*) = \frac{|\lambda_* - \sigma|}{\lambda_* - \sigma} v_* = \pm v_*.$$

In order to show (b), suppose that v_* is such that $\varphi(v_*) = \pm v_*$. Equation (2.3) leads to

$$(J(v_*) - \sigma I)v_* = \frac{\pm 1}{\|(J(v_*) - \sigma I)^{-1}v_*\|} v_*.$$

From Lemma 2.1 it follows that $J(v_*)v_* = A(v_*)v_*$ and we finally have that

$$A(v_*)v_* = \left(\sigma \pm \frac{1}{\|(J(v_*) - \sigma I)^{-1}v_*\|} \right) v_*,$$

and the formula (2.6) follows by multiplying from the left with v_*^T . \square

3. Local convergence properties.

3.1. First-order behavior and convergence factor. From Proposition 2.2 we can directly conclude that if we at some point in the iteration have $v_k = v_*$, where v_* is an eigenvector, then every subsequent iterate v_j , $j > k$, will also be an eigenvector corresponding to the same eigenvalue, but v_j can possibly alternate between $\pm v_*$. We will now study the case where v_k is close but not equal to an eigenvector. In order to understand this local convergence behavior, we also need to take the alternation into account. The error is to first order given by the following result.

THEOREM 3.1 (local convergence). Suppose (λ_*, v_*) is a solution to (1.1) with $\|v_*\| = 1$. Let $\sigma \in \mathbb{R}$ be any shift such that $J(v_*) - \sigma I$ is nonsingular, in particular $\sigma \neq \lambda_*$. Then,

$$(3.1) \quad \varphi'(v_*) = |\lambda_* - \sigma|(I - v_*v_*^T)(J(v_*) - \sigma I)^{-1}.$$

Moreover, suppose the iterates v_k generated by (1.3) are such that $J(v_k) - \sigma I$ is nonsingular for any k . Then, the iterates satisfy

$$(3.2) \quad v_{k+1} \mp v_* = \varphi'(v_*)(v_k - v_*) + O(\|v_k - v_*\|^2).$$

The sign in (3.2) should be chosen negative if $\sigma < \lambda_*$, otherwise positive.

Proof. By using the Taylor expansion and the fixed point characterization in Proposition 2.2, we have that the left-hand side of (3.2) satisfies

$$v_{k+1} \mp v_* = v_{k+1} - (\pm v_*) = \varphi(v_k) - \varphi(v_*) = \varphi'(v_*)(v_k - v_*) + O(\|v_k - v_*\|^2).$$

It remains to derive the formula $\varphi'(v_*)$ given by (3.1).

We will need the derivative of the two-norm, which can be expressed as the row vector

$$(3.3) \quad \frac{\partial}{\partial v} \|\psi(v)\| = \frac{\partial}{\partial v} \sqrt{\psi(v)^T \psi(v)} = \frac{1}{\sqrt{\psi(v)^T \psi(v)}} \psi(v)^T \psi'(v) = \varphi(v)^T \psi'(v),$$

since $\varphi(v)^T = \psi(v)^T / \|\psi(v)\|_2$.

From the definition (2.3) of φ we have the relation $\varphi(v)\|\psi(v)\| = \psi(v)$, whose Jacobian can be computed with the product rule and (3.3),

$$\varphi'(v)\|\psi(v)\| + \varphi(v)\varphi(v)^T \psi'(v) = \psi'(v).$$

Hence,

$$(3.4) \quad \varphi'(v) = \frac{1}{\|\psi(v)\|} (I - \varphi(v)\varphi(v)^T) \psi'(v).$$

Now recall (from Proposition 2.2) that $\varphi(v_*) = \pm v_*$ and consequently $\varphi(v_*)\varphi(v_*)^T = v_* v_*^T$. This fact combined with the formula for the norm (2.7) leads to a simplification of (3.4) when we evaluate at $v = v_*$,

$$(3.5) \quad \varphi'(v_*) = |\lambda_* - \sigma| (I - v_* v_*^T) \psi'(v_*).$$

It remains to establish a formula for the Jacobian of ψ evaluated at $v = v_*$. By differentiation of (2.4) multiplied by $J(v) - \sigma I$, we have

$$(3.6) \quad \left(\frac{\partial}{\partial v} (J(v) - \sigma I) \psi(v) \right)_{\hat{v}=v} + (J(v) - \sigma I) \psi'(v) = I$$

and

$$(3.7) \quad \psi'(v_*) = -(J(v_*) - \sigma I)^{-1} \left(\frac{\partial}{\partial v} J(v) \psi(v_*) \right)_{v=v_*} + (J(v_*) - \sigma I)^{-1}.$$

Moreover, note that $\psi(v_*) = \|\psi(v_*)\| \varphi(v_*) = \pm \|(J(v_*) - \sigma I)^{-1} v_*\| v_*$. We will now show that the first term in (3.7) vanishes identically by showing that all columns of $(\frac{\partial}{\partial v} J(v) v_*)_{v=v_*}$ vanish. Let the j th column be denoted

$$(3.8) \quad c_j := \left(\frac{\partial}{\partial v} J(v) v_* \right)_{v=v_*} e_j = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (J(v_* + \varepsilon e_j) v_* - J(v_*) v_*).$$

Lemma 2.1, and in particular the identity $J(u)u = A(u)u$ for any $u \in \mathbb{R}^n$, implies that

$$\begin{aligned} & J(v_* + \varepsilon e_j)v_* - J(v_*)v_* \\ &= J(v_* + \varepsilon e_j)(v_* + \varepsilon e_j) - \varepsilon J(v_* + \varepsilon e_j)e_j - J(v_*)v_* \\ &= A(v_* + \varepsilon e_j)(v_* + \varepsilon e_j) - \varepsilon J(v_* + \varepsilon e_j)e_j - J(v_*)v_* \\ &= A(v_*)v_* + \varepsilon J(v_*)e_j - \varepsilon J(v_* + \varepsilon e_j)e_j - J(v_*)v_* + O(\varepsilon^2), \end{aligned}$$

where we used a Taylor expansion of $A(v_* + \varepsilon e_j)(v_* + \varepsilon e_j) = A(v_*)v_* + J(v_*)(\varepsilon e_j) + O(\varepsilon^2)$ in the last step. Hence, by again applying Lemma 2.1 now giving $A(v_*)v_* = J(v_*)v_*$, we have

$$c_j = \lim_{\varepsilon \rightarrow 0} (J(v_*)e_j - J(v_* + \varepsilon e_j)e_j + O(\varepsilon)) = 0.$$

We have shown that $c_j = 0$ for any $j = 1, \dots, n$ which implies that

$$(3.9) \quad \left(\frac{\partial}{\partial v} J(v)v_* \right)_{v=v_*} = 0 \in \mathbb{R}^{n \times n}$$

such that the first term in (3.7) vanishes and the proof is complete. \square

Note again that the iteration (1.3) is a fixed point iteration. Hence, the convergence factor is generically given by the spectral radius of the linearized map, i.e., $\varphi'(v_*)$ (if the sign is appropriately taken into account). We will now explicitly characterize the spectral radius of $\varphi'(v_*)$ by deriving formulas for the eigenvalues of $\varphi'(v_*)$. The eigenvalues of $\varphi'(v_*)$ can be explicitly found by multiplying from the left with left eigenvectors of $J(v_*)$. Suppose λ_* is a simple eigenvalue of $J(v_*)$ and suppose $w^T J(v_*) = \mu w^T$, where $\mu \neq \lambda_*$. The fact that a left eigenvector is orthogonal to the right eigenvectors corresponding to all other eigenvalues, we have that $w^T v_* = 0$ and

$$w^T \varphi'(v_*) = |\lambda_* - \sigma| w^T (J(v_*) - \sigma I)^{-1} = \frac{|\lambda_* - \sigma|}{\mu - \sigma} w^T.$$

Under the condition that λ_* is a simple eigenvalue of $J(v_*)$, this provides an explicit formula for $n - 1$ eigenvalues of $J(v_*)$. The final eigenvalue is 0 and the largest eigenvalue in modulus, i.e., the spectral radius, is

$$(3.10) \quad \gamma := \rho(\varphi(v_*)) = \frac{|\lambda_* - \sigma|}{|\mu_2 - \sigma|},$$

where $\mu_2 \in \mathbb{C}$ is the eigenvalue of $J(v_*)$ closest to σ , but not equal to λ_* , i.e.,

$$\mu_2 = \operatorname{argmin}_{\mu \in \lambda(J(v_*)) \setminus \{\lambda_*\}} |\mu - \sigma|.$$

Consequently, a sequence $\{v_k\}_{k=0}^\infty$ that converges to v_* satisfies

$$(3.11) \quad \lim_{k \rightarrow \infty} \frac{\|v_{k+1} \pm v_*\|}{\|v_k \pm v_*\|} = \gamma$$

under generic conditions. A formalization of this statement can be done with the concepts of R-convergence and Q-convergence in [25, NR 10.1-5].

3.2. Similarities and differences with inverse iteration for linear problems. The iteration (1.3) is clearly a generalization of inverse iteration when $A(v) = A_0$ is constant, and a behavior similar to (linear) inverse iteration is expected when the problem is close to linear. More importantly, from the theory above we can conclude that the iteration possesses many properties similar to inverse iteration also in a situation when the problem is not close to linear.

- The convergence is in general linear.
- The convergence factor is asymptotically proportional to eigenvalue shift distance when the distance is small.
- The convergence factor (3.10) is in general determined by a quotient involving the eigenvalue shift distance in the numerator. Unlike the linear case, the denominator is the distance between the shift and the second closest eigenvalue of $J(v_*)$.
- Unlike the linear case, for a given shift, several eigenvectors can be attractive fixed points. Conversely, for a given shift, the iteration may not have any attraction points.

In order to further illustrate the value of these properties we will in this work also briefly study another generalization of inverse iteration. Inverse iteration for standard eigenvalue problems consists of shifting, inverting, and normalizing, and it is from this perspective natural to consider the generalization $v_{k+1} = \varphi_A(v_k)$,

$$(3.12) \quad \varphi_A(v) := \frac{1}{\|\psi_A(v)\|} \psi_A(v), \quad \psi_A(v) = (A(v) - \sigma I)^{-1}v.$$

We will call this the A -version of inverse iteration.

We can derive a description of the first-order behavior similar to Theorem 3.1. Note that $\varphi'_A(v_*)$ is considerably more complicated than $\varphi'(v_*)$ and most of the bullets above do not apply to this version, most importantly, the convergence factor is not necessarily small if the shift is close to the eigenvalue.

3.3. Illustration of local convergence. Several properties of the algorithm, including the local convergence above, can be observed when applied to the example

$$A(v) = A_0 + \sin\left(\frac{v^T B v}{v^T v}\right) A_1.$$

The Jacobian is given by

$$J(v) = \frac{\partial}{\partial v} A(v)v = A(v) + 2 \frac{\cos\left(\frac{v^T B v}{v^T v}\right)}{(v^T v)^2} A_1 v ((v^T v)v^T B - (v^T B v)v^T).$$

We selected A_0 , A_1 , and B in a random way. In order to make the numerical simulations reproducible, we will fix A_0 , A_1 , and B as follows:

$$A_0 = \frac{1}{10} \begin{pmatrix} 10 & 21 & 13 & 16 \\ 21 & -26 & 24 & 2 \\ 13 & 24 & -26 & 37 \\ 16 & 2 & 37 & -4 \end{pmatrix}, \quad A_1 = \frac{\beta}{10} \begin{pmatrix} 20 & 28 & 12 & 32 \\ 28 & 4 & 14 & 6 \\ 12 & 14 & 32 & 34 \\ 32 & 6 & 34 & 16 \end{pmatrix},$$

$$B = \frac{1}{10} \begin{pmatrix} -14 & 16 & -4 & 15 \\ 16 & 10 & 15 & -9 \\ -4 & 15 & 16 & 6 \\ 15 & -9 & 6 & -6 \end{pmatrix},$$

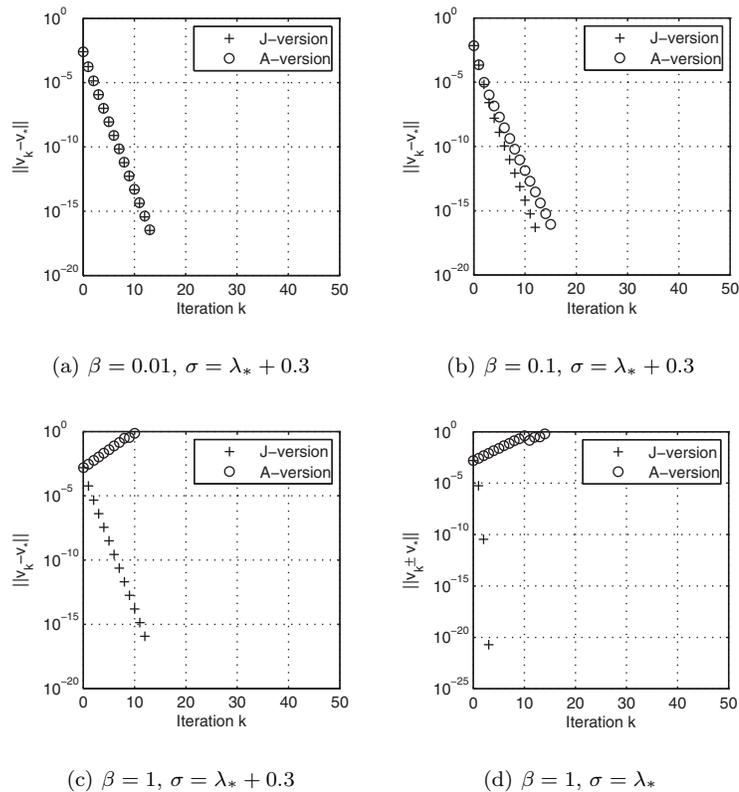


FIG. 1. Convergence for the example in section 3.3 using (1.3), i.e., the J-version of inverse iteration, and the A-version of inverse iteration (3.12).

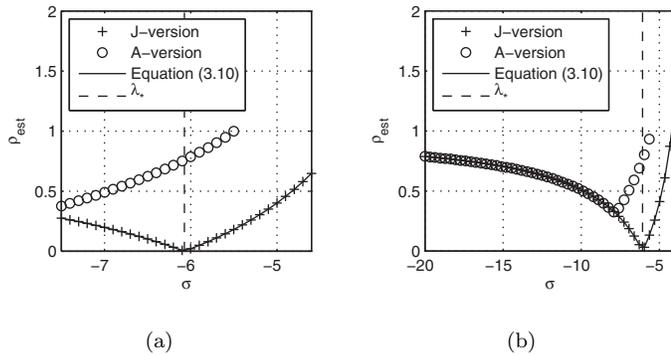


FIG. 2. Estimated convergence factor for the J-version and A-version of iteration, where estimation is done with the quotient $\|v_{k+1} \pm v_K\|/\|v_k \pm v_K\|$, where v_K is a very accurate solution. Clearly, when $\sigma \rightarrow \lambda_*$ the convergence factor for the J-version approaches zero, unlike the A-version. When $\sigma \rightarrow -\infty$ the convergence factors coincide. We have selected $\beta = 0.5$.

where we carry out simulations for a number of different β . Note that the problem is linear when $\beta = 0$.

Simulations of the inverse iterations for different β are given in Figures 1 and 2. In Figure 1 we clearly observe linear convergence, for the A-version as well as the J-version. Moreover, for small β , i.e., when the problem is close to linear, we see

the performance of the A -version is essentially as good as for the J -version. The J -version is clearly less dependent on β . As the nonlinearity parameter β increases, the performance of the A -version worsens whereas the convergence of the J -version is not greatly affected by an increase of β . We see in Figure 1(d) that the slow convergence of the A -version can not always be compensated for by moving σ closer to the eigenvalue λ_* . In this case the A -version is not even locally convergent when $\sigma = \lambda_*$, whereas the J -version exhibits quadratic convergence. Note that for $\sigma = \lambda_*$ the fixed point maps $\varphi(v_*)$ and $\varphi_A(v_*)$ involve inverses of singular matrices, implying that they do formally not have fixed points v_* . Similar to inverse iteration for standard eigenvalue problems the J -version still works in practice when subject to rounding errors.

In Figure 2 we visualize the convergence factor predicted by formula (3.10) with an experimental convergence factor determined by starting the iteration very close to the eigenvector. We see that the convergence factor of the J -version approaches zero when $\sigma \rightarrow \lambda_*$ as predicted by (3.11). The convergence factor for the A -version does clearly not vanish when σ is close to λ_* .

3.4. Updating shift and connection with Newton's method. A standard approach to accelerate local convergence of inverse-iteration-type methods consists of updating the shift, based on the eigenvector approximation. In what follows we derive an updating strategy, having an interpretation in terms of Newton's method.

Given a vector $\ell \in \mathbb{R}^n$, an application of Newton's methods to the following system of equations,

$$(3.13) \quad \begin{cases} (A(v) - \lambda I)v = 0, \\ \ell^T v - 1 = 0, \end{cases}$$

leads to the iteration

$$\begin{bmatrix} J(v_k) - \lambda_k I & -v_k \\ \ell^T & 0 \end{bmatrix} \begin{bmatrix} v_{k+1} - v_k \\ \lambda_{k+1} - \lambda_k \end{bmatrix} = - \begin{bmatrix} (A(v_k) - \lambda_k I)v_k \\ \ell^T v_k - 1 \end{bmatrix}.$$

By using Lemma 2.1 we find that

$$\begin{cases} v_{k+1} = (\lambda_{k+1} - \lambda_k)(J(v_k) - \lambda_k I)^{-1}v_k, \\ \ell^T v_{k+1} = 1. \end{cases}$$

Consequently, the iteration

$$(3.14) \quad v_{k+1} = \beta_k (J(v_k) - \sigma_k I)^{-1}v_k$$

with $\beta_k = 1/\ell^T (J(v_k) - \sigma_k I)^{-1}v_k$ and the update law

$$\sigma_{k+1} = \sigma_k + \beta_k,$$

is equivalent to Newton's method applied to (3.13). As a consequence, quadratic convergence is guaranteed for the case of a simple solution. Note that due to the scale invariance of J , i.e., $J(\alpha v) = J(v)$, inherited from the scale invariance of A , iteration (3.14) is essentially the same as iteration (1.1). The change in normalization allows us to express the update of the shift in a simple way. We also note that as an alternative update strategy, the Rayleigh quotient can be used, i.e., $\sigma_{k+1} = p(v_{k+1})$, with p given by (1.7).

4. Interpretation as the discretization of an ODE.

4.1. The normalized ODE associated with (1.1). In order to provide further insight into the nonlocal behavior of the iteration, we will consider the following ODE:

$$(4.1) \quad z'(t) = -A(z(t))z(t)$$

and the function corresponding to the normalization of the solution

$$(4.2) \quad y(t) := q(t)z(t), \text{ where } q(t) := \frac{1}{\|z(t)\|}.$$

The function y also satisfies an ODE, which does not involve z . This can be seen from the following reasoning. Note that

$$q'(t) = \frac{-1}{2(z(t)^T z(t))^{3/2}} 2z(t)^T z'(t) = \frac{1}{\|z(t)\|^3} z(t)^T A(z(t))z(t) = q(t)y(t)^T A(y(t))y(t).$$

By applying the product rule to (4.2), we have

$$y'(t) = q(t)z'(t) + q'(t)z(t) = -q(t)A(z(t))z(t) + q(t)(y(t)^T A(y(t))y(t))z(t).$$

Hence, the normalized function $y(t)$ satisfies the differential equation

$$(4.3) \quad y'(t) = p(y(t))y(t) - A(y(t))y(t),$$

where $p(y(t))$ is the *Rayleigh quotient* defined by (1.7). The ODE (4.3) will be used as a tool to understand the iteration (1.3). For this reason several properties of the ODE and some of its relations with the nonlinear eigenvalue problem (1.1) will be particularly important.

THEOREM 4.1 (ODE properties). *Consider the ODE (4.3) with initial condition $y(0) = y_0$ with $\|y_0\| = 1$. The ODE and its solution have the following properties.*

- (a) *The norm $\|y(t)\| = 1$ independent of t .*
- (b) *Any stationary point of the ODE (4.3) is a normalized eigenvector of (1.1).*
- (c) *Any normalized eigenvector of (1.1) is a stationary point of the ODE (4.3).*
- (d) *Let y_* be a stationary point of the ODE (4.3). The stationary point is stable if the eigenvalue $\lambda_* = p(y_*)$ is a simple dominant eigenvalue of $J(y_*)$, i.e., if λ_* is a simple eigenvalue of $J(y_*)$ and*

$$\lambda_* < \operatorname{Re}(\mu) \text{ for any } \mu \in \lambda(J(y_*)) \setminus \{\lambda_*\}.$$

- (e) *Let y_* be a stationary point of the ODE (4.3). The stationary point is unstable if there is an eigenvalue μ of $J(y_*)$ such that*

$$\operatorname{Re}(\mu) < \lambda_* = p(y_*).$$

Proof. The statements (a)–(c) follow directly from the derivation of (4.3). In order to study the stability of the stationary point y_* , let $\Delta(t) := y(t) - y_*$ denote the deviation from the stationary solution. Then, by Taylor expansion of (4.3) we have

$$\Delta'(t) = y'(t) = p(y(t))y(t) - A(y(t))y(t) = [p(y_*)I + y_* p'(y_*) - J(y_*)] \Delta(t) + O(\Delta(t))^2.$$

We can now insert $p(y_*) = \lambda_*$ and $p'(y_*) = y_*^T(J(y_*) - \lambda_*I)$ since $\|y_*\| = \|y_0\| = 1$. Hence, the behavior is, to first order, given by

$$(4.4) \quad \tilde{\Delta}'(t) = F(y_*)\tilde{\Delta}(t),$$

where

$$F(y_*) = \lambda_*I + y_*y_*^T(J(y_*) - \lambda_*I) - J(y_*) = (I - y_*y_*^T)(\lambda_*I - J(y_*)).$$

Since $F(y_*)$ is singular, at least one eigenvalue is zero. The other eigenvalues can be determined by multiplying $F(y_*)$ from the left with left eigenvectors of w . Suppose w is a left eigenvector corresponding to any eigenvalue $\mu \neq \lambda_*$ of $J(y_*)$. Due to the fact that a left eigenvector is orthogonal to the right eigenvectors corresponding to other eigenvalues we have that $w^T y_* = 0$ and $w^T F(y_*) = w^T(I - y_*y_*^T)(\lambda_*I - J(y_*)) = w^T(\lambda_*I - J(y_*)) = (\lambda_* - \mu)w^T$. Hence, all of the nonzero eigenvalues of $F(y_*)$ have the form $\lambda_* - \mu$, where μ is an eigenvalue of $J(y_*)$, but not equal to λ_* . The ODE (4.4) is unstable if $F(y_*)$ has eigenvalues with positive real part, and the ODE (4.4) is stable if all eigenvalues have negative real part. The statements (d)–(e) follow from the sign of $\text{Re}(\lambda_* - \mu) = \lambda_* - \text{Re}(\mu)$, which, provided it is not zero, gives a conclusion about the local stability of (4.4) and (4.3). \square

4.2. Discretization of the ODE and equivalence with inverse iteration.

In the previous subsection we saw how the ODE (4.3) was related to the nonlinear eigenvalue problem (1.1); in particular we showed that the stationary solutions were equivalent to the eigenvectors of (1.1). We will now show how this connection can be used by showing that the proposed version of inverse iteration (1.3) can be interpreted as a discretization of the ODE, allowing us to study the general behavior of the iteration by studying the ODE.

Consider first the Rosenbrock–Euler method [12, Chapter IV.7], i.e., the backward Euler method where the nonlinear system is solved with one step of Newton–Raphson’s method. The backward Euler method with step length h applied to (4.3) is

$$(4.5) \quad \frac{\tilde{y}_{k+1} - \tilde{y}_k}{h} \approx p(\tilde{y}_{k+1})\tilde{y}_{k+1} - A(\tilde{y}_{k+1})\tilde{y}_{k+1}.$$

We approximate the first term in (4.5) by its linearization,

$$p(\tilde{y}_{k+1})\tilde{y}_{k+1} \approx p(\tilde{y}_k)\tilde{y}_k + (p'(\tilde{y}_k)\tilde{y}_k + p(\tilde{y}_k))(\tilde{y}_{k+1} - \tilde{y}_k) = p(\tilde{y}_k)\tilde{y}_{k+1},$$

where we denoted $p'(z) := \frac{\partial}{\partial z}p(z) \in \mathbb{R}^{1 \times n}$. We also used that $p'(z)z = 0$ for any z which follows from Lemma 2.1 since $p(z)$ is scaling invariant. The second term in (4.5) is now also approximated by its linearization. By combining the approximation with (2.2), we have

$$A(\tilde{y}_{k+1})\tilde{y}_{k+1} \approx A(\tilde{y}_k)\tilde{y}_k + J(\tilde{y}_k)(\tilde{y}_{k+1} - \tilde{y}_k) = J(\tilde{y}_k)\tilde{y}_{k+1}.$$

Hence, the Rosenbrock–Euler method applied to (4.3) is equivalent to

$$\frac{1}{h}(\tilde{y}_{k+1} - \tilde{y}_k) = p(\tilde{y}_k)\tilde{y}_{k+1} - J(\tilde{y}_k)\tilde{y}_{k+1}$$

and also

$$(4.6) \quad \left(\left(\frac{1}{h} - p(\tilde{y}_k) \right) I + J(\tilde{y}_k) \right) \tilde{y}_{k+1} = \frac{1}{h}\tilde{y}_k.$$

The ODE (4.3) has, according to Theorem 4.1, a solution with constant norm. Despite this, the discretization \tilde{y}_k does not have constant norm, due to the approximation error introduced in the stepping scheme. This issue can be addressed by carrying out a *standard projection* described, e.g., in [11, Algorithm IV.4.2], which is a common procedure for incorporating normalization constraints. In this case, it reduces to subsequently normalizing the iterate, i.e., setting $y_{k+1} := \frac{1}{\|\tilde{y}_{k+1}\|} \tilde{y}_{k+1}$. Since p and J are scaling invariant, we have $p(\tilde{y}_k) = p(y_k)$ and $J(\tilde{y}_k) = J(y_k)$, and we can express one step of the integration scheme as

$$(4.7) \quad y_{k+1} = \alpha_k \left(\left(\frac{1}{h} - p(y_k) \right) I + J(y_k) \right)^{-1} y_k,$$

where $\alpha_k = 1 / \left\| \left(\frac{1}{h} - p(y_k) \right) I + J(y_k) \right\|^{-1} y_k$.

The main result in this section is based on the observation that (4.7) is equivalent to the proposed version of inverse iteration (1.3) if the shift σ and step length h are related in a particular way. This connection is summarized in the following theorem.

THEOREM 4.2 (ODE discretization equivalence). *Let the sequences $\{y_k\}_{k=0}^{\infty}$ and $\{v_k\}_{k=0}^{\infty}$ be generated in the following way.*

- (a) *Let v_k be the iterates generated by inverse iteration with (1.3) with a given shift $\sigma \in \mathbb{R}$.*
- (b) *Let y_k be the discretization of (4.3) using*
 - *the Rosenbrock–Euler method,*
 - *a standard projection step which imposes the normalization, and*
 - *the step length*

$$(4.8) \quad h_k = \frac{1}{p(y_k) - \sigma},$$

where σ is the shift used to generate v_k .

Moreover, suppose they are started in the same way, i.e., let $y_0 = v_0$. Then,

$$v_k = y_k$$

for all $k \in \mathbb{N}$.

We summarize some immediate consequences.

- The iteration (1.3) is an approximation of the trajectory of the ODE (4.3) if h is chosen sufficiently small, i.e., if σ is chosen sufficiently negative.
- Suppose the ODE converges to a stationary solution, which indeed is the case for several examples, and can in particular be proven for the example in section 5. In this situation the stationary solution is a solution to the ODE, and the iteration will converge to an eigenvector for sufficiently negative σ .
- For the linear case, i.e., $A(v) = B$, we can explicitly write down the solution to (4.1) using the Jordan form, and thereby study (4.3) which will converge to the dominant eigenvalue unless $y(0)$ is chosen in a particular way. Moreover, Theorem 4.1(d)–(e) shows that if the dominant eigenvalue is simple, it will be the only stable solution. Analogously, for nonlinear problems which resemble the linear case in the sense that the ODE (4.3) only has one stable stationary solution and the ODE is convergent (unless started in a particular way), the iteration will converge to that solution for sufficiently negative σ .

Remark 4.3 (ODE interpretation for A-version). The A-version (3.12) can be interpreted as a discretization of the ODE (4.3), but where J is approximated by

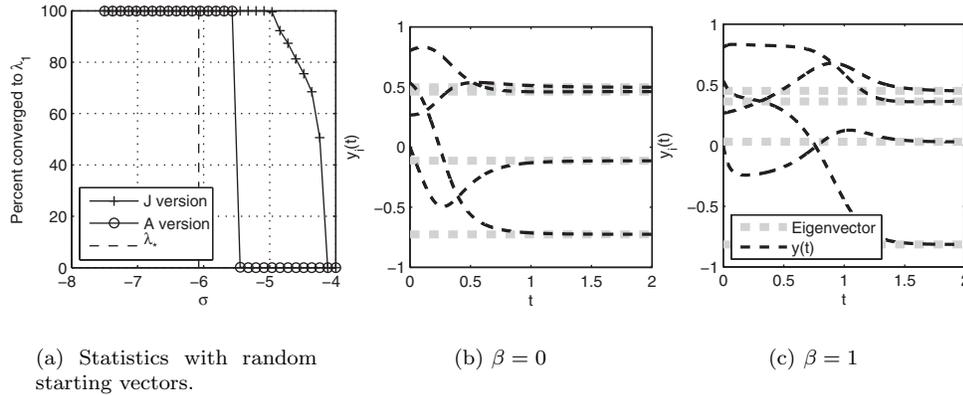


FIG. 3. Convergence of the ODE (sine example) to the dominant eigenvector.

A and h is chosen according to (4.8). In fact, when the A -version is applied to the example in section 5 it is equivalent to an approach available in the literature; more precisely, the discretization of the normalized gradient flow given in [1, equations (2.20)–(2.21)].

4.3. Illustration and numerical justification of ODE discretization. Consider a situation where we do not have any approximation of any eigenvalue of the example in section 3.3. To characterize this situation we carry out a number of simulations with random starting vectors for different shifts. Such a statistical simulation is shown in Figure 3a. We observe that for $\sigma < -5$, the iteration appears to always converge to the dominant solution. This property can be explained using the developed theory. When $\sigma \ll \lambda_*$ we can apply Theorem 4.2 and when $|\sigma - \lambda_*| \ll 1$ we can apply the local convergence theorem (Theorem 3.1).

We computed the solution for the ODE (4.3) for this example using a standard ODE integrator, with a random starting value. The trajectory is shown in Figures 3(b)–(c). Clearly, the ODE converges to the dominant eigenvector. Similarly, in Figure 4 we have shown the computed approximate trajectories using (1.3) (J -version) and (3.12) (A -version). We see in accordance with the interpretation of σ as a particular choice of step length (via (4.8)) that the iteration follows the trajectories better if σ is more negative. Moreover, with the J -version we can take larger steps than with the A -version, as the A -version does not follow the ODE in Figure 4(a).

The experience with the ODE (1.3) for this particular example indicates that the ODE converges to a stationary solution for all $\beta \in [0, 1]$. Moreover, for several choices of β , among all eigenvectors, there appears to be only one situation where λ_* is the dominant eigenvalue of $J(v_*)$. This can be seen in Figure 5. From Theorem 4.1 we consequently know that the ODE (1.3) only has one stable stationary solution. Combined with the convergence observation, the ODE will always converge to this solution. Hence, Theorem 4.2 implies that for sufficiently negative σ , the iteration will follow the trajectory of the ODE and eventually converge to the dominant eigenvector.

5. Application to the Gross–Pitaevskii equation.

5.1. Discretization and reformulation to the form (1.1). To achieve the physical phenomenon of Bose–Einstein condensation, weakly interacting particles (such as ^4He atoms) are trapped and cooled down to very low temperatures, thereby undergoing a phase transition into a *Bose–Einstein condensate* (BEC), a superfluid phase

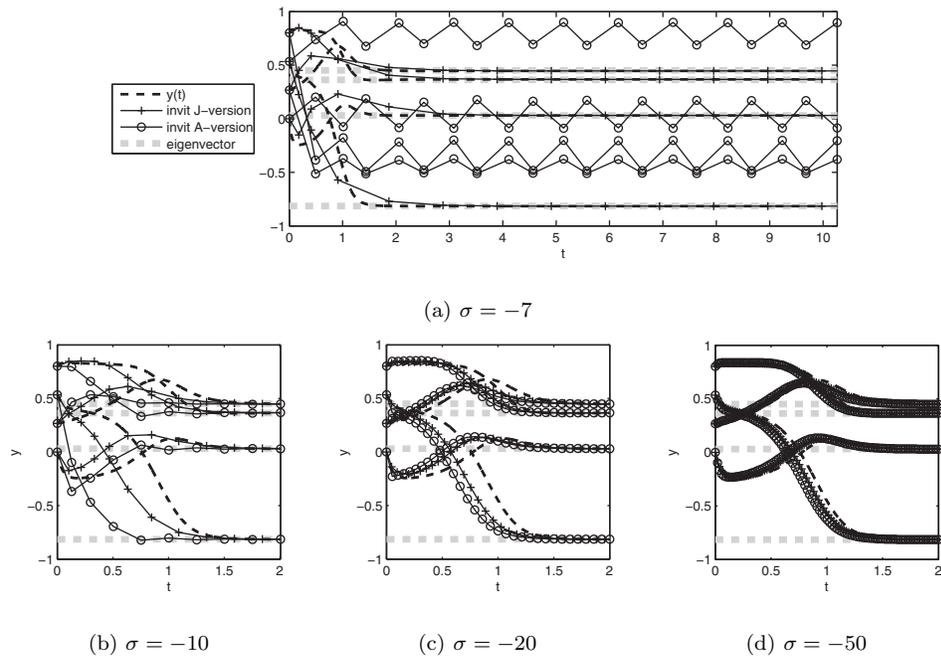


FIG. 4. The *J*-version is better in the sense that it follows the ODE also for larger h than for the *A*-version. The dominant eigenvalue is $\lambda_* \approx -6.01$.

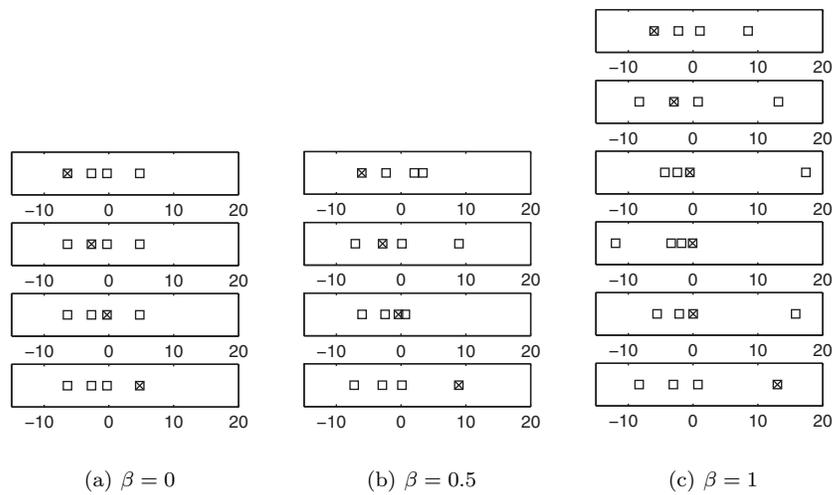


FIG. 5. The eigenvalues of $J(v_*)$ for the example discussed in section 4.3. In every subfigure, \square denotes the eigenvalues of $J(v_*)$ for an eigenpair (λ_*, v_*) of (1.1) and λ_* is denoted by \times . In this notation, a stable stationary solution described by Theorem 4.1(d)–(e) corresponds to a situation where \times is to the left of all \square .

of collective quantum mechanical motion. A standard model for a BEC is the GPE, a nonlinear partial differential eigenvalue equation. This equation exhibits interesting topological behavior. In particular, it reproduces experimentally observed quantized vortices in a BEC within a rotating trap, the hallmark of superfluidity.

Downloaded 03/17/15 to 134.58.46.142. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Upon discretization, the GPE becomes an eigenvalue problem of the form (1.1), and to illustrate the iteration (1.3) we will consider the case of a rotating BEC on the domain $\mathbb{D} = (-L, L) \times (-L, L)$ in two dimensions, whose GPE reads

$$(5.1) \quad \left(-\frac{1}{2}\Delta - i\Omega \frac{\partial}{\partial \phi} + V(x, y) \right) \psi(x, y) + b|\psi(x, y)|^2 \psi(x, y) = \lambda \psi(x, y), \quad (x, y) \in \mathbb{D},$$

where $\psi(x, y)$ is the condensate wave function (not to be confused with ψ in (2.4)) such that $|\psi(x, y)|^2$ is proportional to the *particle density* at (x, y) , and where $\frac{\partial}{\partial \phi}$ is the angular derivative given by

$$(5.2) \quad \frac{\partial}{\partial \phi} = y \frac{\partial}{\partial x} - x \frac{\partial}{\partial y}.$$

The function $V(x, y)$ is the external potential function which describes the shape of the particle trap. Here, we choose an asymmetric harmonic oscillator potential $V(x, y) = (x^2 + 1.2y^2)/2$. We are mostly interested in those solutions to (5.1) which are physically important, e.g., the ground state.

The boundary condition is chosen as $\psi(x, y) = 0$ on $(x, y) \in \partial\mathbb{D}$. Moreover, ψ is normalized such that $\|\psi\|_{L^2(\mathbb{D})} = 1$. The physical constant b controls the strength of the interactions between the bosons, and Ω the angular velocity of the rotation, which are here selected as $b = 200$ and $\Omega = 0.85$, respectively.

The domain is discretized with $N + 2$ equidistant grid points in each physical direction leading to $n_c = N^2$ interior grid points with spacing Δx in both directions. For completeness we provide the matrices resulting from the discretization. We use the approximations of Δ and $\partial/\partial\phi$,

$$L_N = D_{2,N} \otimes I + I \otimes D_{2,N}, \quad L_{\phi,N} = \text{diag}(y_1, \dots, y_N) \otimes D_N - D_N \otimes \text{diag}(x_1, \dots, x_N),$$

where D_N and $D_{2,N}$ are the central difference approximations of the derivative and the second derivative, respectively. Moreover, we let

$$\tilde{V} = (V(x_1, y_1), \dots, V(x_N, y_1), V(x_1, y_2), \dots, V(x_N, y_2), \dots, V(x_N, y_N)) \in \mathbb{R}^n.$$

Then, with

$$\tilde{A}_0 = -\frac{1}{2}L_N - i\Omega L_{\phi,N} + \text{diag}(\tilde{V})$$

the discretized problem is

$$(5.3) \quad \tilde{A}(z)z = \lambda z,$$

where

$$(5.4) \quad \tilde{A}(z) = \tilde{A}_0 + \beta \text{diag}(|z|)^2,$$

and to be consistent with $\|\psi\|_{L^2(\mathbb{D})} = 1$ we must have $\psi(x_j, y_k) \approx (\Delta x)^{-1} z_{N(k-1)+j}$ and $\beta = (\Delta x)^{-2} b$.

The problem (5.3) is a complex problem of dimension n_c , not satisfying the scaling invariance (1.2). In order to transform this problem to the form (1.1) we introduce the following notation:

$$\begin{aligned} A(v) &:= \begin{pmatrix} \operatorname{Re} \tilde{A} \left(\frac{v_1 + iv_2}{\sqrt{v_1^T v_1 + v_2^T v_2}} \right) & -\operatorname{Im} \tilde{A} \left(\frac{v_1 + iv_2}{\sqrt{v_1^T v_1 + v_2^T v_2}} \right) \\ \operatorname{Im} \tilde{A} \left(\frac{v_1 + iv_2}{\sqrt{v_1^T v_1 + v_2^T v_2}} \right) & \operatorname{Re} \tilde{A} \left(\frac{v_1 + iv_2}{\sqrt{v_1^T v_1 + v_2^T v_2}} \right) \end{pmatrix} \\ &= \begin{pmatrix} \operatorname{Re} \tilde{A}_0 & -\operatorname{Im} \tilde{A}_0 \\ \operatorname{Im} \tilde{A}_0 & \operatorname{Re} \tilde{A}_0 \end{pmatrix} + \frac{\beta}{v^T v} B(v), \end{aligned}$$

and

$$B(v) := \begin{pmatrix} \operatorname{diag}(v_1)^2 + \operatorname{diag}(v_2)^2 & 0 \\ 0 & \operatorname{diag}(v_1)^2 + \operatorname{diag}(v_2)^2 \end{pmatrix}, \quad v = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix}.$$

With this definition of the matrix $A(v)$ we have transformed $\tilde{A}(z)$ by treating the real and imaginary parts of z as separate variables. Consequently, the complex problem (5.3) of dimension n_c is equivalent to the real problem

$$(5.5) \quad A(v)v = \lambda v$$

of dimension $n = 2n_c$ which does satisfy (1.2).

5.2. Jacobian and exploitation of Jacobian structure. The resulting matrix $A(v)$ in (5.5) is a sparse matrix for any v due to the finite-difference discretization. In order to carry out (1.3) we need the Jacobian $J(v)$ and we need to be able to efficiently solve the corresponding shifted linear system of equations. It turns out that the Jacobian is in general a full matrix, making the direct application of sparse solvers inefficient. Fortunately, the linear system involving the Jacobian can be decomposed into two linear systems of equations involving a matrix which is sparse. The derivation is based on the fact that the Jacobian is the sum of a sparse matrix and a rank-one matrix. For such structures the Sherman–Morrison–Woodbury formula [10, section 2.1.3] is a standard technique.

THEOREM 5.1 (the Jacobian associated with the GPE). *The Jacobian for the nonlinear eigenvalue problem (5.5) corresponding to the GPE is given by*

$$(5.6) \quad \begin{aligned} J(v) &:= \frac{\partial}{\partial v} (A(v)v) = \begin{pmatrix} \operatorname{Re} A_0 & -\operatorname{Im} A_0 \\ \operatorname{Im} A_0 & \operatorname{Re} A_0 \end{pmatrix} \\ &\quad + \frac{\beta}{v^T v} \left[\begin{pmatrix} 3 \operatorname{diag}(v_1)^2 + \operatorname{diag}(v_2)^2 & 2 \operatorname{diag}(v_1) \operatorname{diag}(v_2) \\ 2 \operatorname{diag}(v_1) \operatorname{diag}(v_2) & \operatorname{diag}(v_1)^2 + 3 \operatorname{diag}(v_2)^2 \end{pmatrix} \right. \\ &\quad \left. - \frac{2}{v^T v} B(v)vv^T \right], \end{aligned}$$

where $v^T = (v_1^T, v_2^T)$ with $v_1, v_2 \in \mathbb{R}^{n/2}$. Moreover, suppose $v^T v = 1$ and let

$$\begin{aligned} C &:= \begin{pmatrix} \operatorname{Re} A_0 & -\operatorname{Im} A_0 \\ \operatorname{Im} A_0 & \operatorname{Re} A_0 \end{pmatrix} \\ &\quad + \beta \begin{pmatrix} 3 \operatorname{diag}(v_1)^2 + \operatorname{diag}(v_2)^2 & 2 \operatorname{diag}(v_1) \operatorname{diag}(v_2) \\ 2 \operatorname{diag}(v_1) \operatorname{diag}(v_2) & \operatorname{diag}(v_1)^2 + 3 \operatorname{diag}(v_2)^2 \end{pmatrix} - \sigma I \end{aligned}$$

and suppose C is nonsingular. Then, the solution to the linear system $(J(v) - \sigma I)^{-1}v$ can be expressed as

$$(5.7) \quad (J(v) - \sigma I)^{-1}v = u_1 + u_2,$$

where

$$u_1 = C^{-1}v$$

and

$$u_2 = \frac{v^T u_1}{1 - v^T w} w \quad \text{with} \quad w = 2\beta C^{-1}B(v)v.$$

Proof. The formula for the Jacobian (5.6) follows from several applications of the product rule. More precisely, we have

$$\begin{aligned} \frac{\partial}{\partial v} A(v)v &= \frac{\partial}{\partial v} \left(\begin{pmatrix} \operatorname{Re} A_0 & -\operatorname{Im} A_0 \\ \operatorname{Im} A_0 & \operatorname{Re} A_0 \end{pmatrix} v + \frac{\beta}{v^T v} B(v)v \right) \\ &= \begin{pmatrix} \operatorname{Re} A_0 & -\operatorname{Im} A_0 \\ \operatorname{Im} A_0 & \operatorname{Re} A_0 \end{pmatrix} + \frac{\beta}{(v^T v)^2} \left(v^T v \left(\frac{\partial}{\partial v} B(v)v \right) - 2B(v)vv^T \right). \end{aligned}$$

The term involving the Jacobian of $B(v)v$ can now be simplified,

$$\begin{aligned} \frac{\partial}{\partial v} B(v)v &= \frac{\partial}{\partial v} \left(\operatorname{diag}(v)^2 v + \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \operatorname{diag}(v)^2 \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} v \right) \\ &= 3 \operatorname{diag}(v)^2 + \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \operatorname{diag}(v)^2 \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \\ &\quad + 2 \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \operatorname{diag}(v) \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \operatorname{diag}(v) \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \\ &= B(v) + 2 \operatorname{diag}(v)^2 + 2 \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \operatorname{diag}(v) \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix} \operatorname{diag}(v) \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix}. \end{aligned}$$

The relation (5.6) follows from expanding and combining the three terms.

The formula for (5.7) follows from the fact that the last term in (5.6) is a rank-one term and we can apply the Sherman–Morrison–Woodbury formula [10, section 2.1.3]. \square

5.3. Specialized ODE interpretation and step-length heuristics. We will now use the following important observation. The ODE (4.3) is equivalent to the ODE representing a flow in the technique known as imaginary time propagation or normalized gradient flow in, e.g., [1]. In particular, (4.3) is equivalent to [1, equations (2.15)–(2.16)]. This connection allows us to directly reach conclusions about the ODE (4.3) for the GPE. We conclude from [1, Remark 2.6] that the ODE will converge to a stationary solution. It can equivalently be shown to converge by constructing a Lyapunov function and applying a variant of Lyapunov’s second method (see [13, Chapter 3] for Lyapunov’s second method). Moreover, the imaginary time propagation technique is for physical reasons known to converge to a physically relevant solution (there may be several), e.g., the ground state or metastable configuration of the system.

In light of the equivalence with imaginary time propagation, it is natural to follow the true flow as closely as possible during the iteration, while at the same time taking

time steps as long as possible to minimize computation time. Therefore, we propose the following heuristic for the step length h or equivalently a choice of σ .

Suppose that we have computed an approximation $y_k \approx y(t_k)$, and we now need to determine an appropriate h to compute $y_{k+1} \approx y(t_k + h)$. A step-length choice for the Rosenbrock–Euler method is given in Appendix A and, in particular, formula (A.2), for a given fixed local error ε . Note that for our case, i.e., the ODE (4.3), $f(v) = p(v)v - A(v)v$ and $f'(v) = -(I - vv^T)(J(v) - p(v)I) + vv^T(A(v) - p(v)I)$ such that

$$f'(v)f(v) = (I - vv^T)(-J(v)f(v) + p(v)f(v)) + vv^T(A(v) - p(v)I)f(v).$$

We will now see that this expression can be efficiently computed before each step of the iterative method. This leads us to the choice of the shift σ we propose to use in this work:

- Compute $f_k := f(v_k) = p(v_k)v_k - A(v_k)y_k$;
- compute $g_k := J(v_k)f_k$;
- compute $e_k := (I - v_k v_k^T)(-e_k + p(v_k)f_k) + v_k v_k^T(A(v_k) - p(v_k)I)f_k$;
- compute $h_k := (2\varepsilon/\|e_k\|)^{1/2}$ using a desired tolerance ε ;
- if $h_k > h_{\max}$ set $h_k = h_{\max}$;
- set $\sigma = p(v_k) - 1/h_k$ according to (4.8).

The choice to make sure that $h_k \leq h_{\max}$ is to avoid taking too large steps, for which the reasoning for step length above is not supported.

Note that we do not need form the matrices $I - v_k v_k^T$ or $J(v_k)$ explicitly when we apply it to the GPE. It is more efficient to instead compute the vectors f_k , g_k , and e_k by forming products between vectors and matrix vector multiplications with sparse matrices since, e.g., $J(v_k)$ is given as the sum of a matrix and a rank-one matrix in Theorem 5.1.

5.4. Conclusions from computational results. We carried out the inverse iteration algorithm with the heuristic choice of σ proposed in section 5.3, for a number of different choices of parameters. We selected $b = 200$, $L = 15$, and $\Omega = 0.85$. The number of grid points is $N = 300$, i.e., the eigenvalue problem (1.1) is of size $n = 180000$. The step-length heuristic was chosen with parameters $\varepsilon = 2$ and $h_{\max} = 10^4$ (except where otherwise stated). An initial guess was chosen as a random superposition of Gaussians, such that its length scale is independent of the interior grid size N . The simulation was completed in $1.2 \cdot 10^3$ seconds with an implementation of the algorithm in MATLAB running on an Apple MacBook Pro with a 2.6 GHz Intel i7 quad-core processor.

The results of the numerical simulations are presented in Figures 6–10. In these figures ψ_k denotes the approximate solution to (5.1) after k iterations, and K denotes the total number of iterations.

The convergence is visualized in Figure 7. As expected from the ODE interpretation in section 4 and the fact that the GPE ODE converges, we eventually reach a solution. Moreover, the asymptotic convergence is fast as the step length is larger when the solution is accurate. The approximations of the solution at different iterates are given in Figure 8, showing the shape of the function as it evolves and converges. Clearly, the random initial condition turns into a physically meaningful approximation after only a few iterations, and the final iterations mostly change the position of the vortices.

We can also observe the local convergence properties presented in section 3 in this application. Note first that when $h = h_{\max}$, we can see that σ is almost constant. This can be observed in Figure 10. It is also expected from the fact that the error behaves

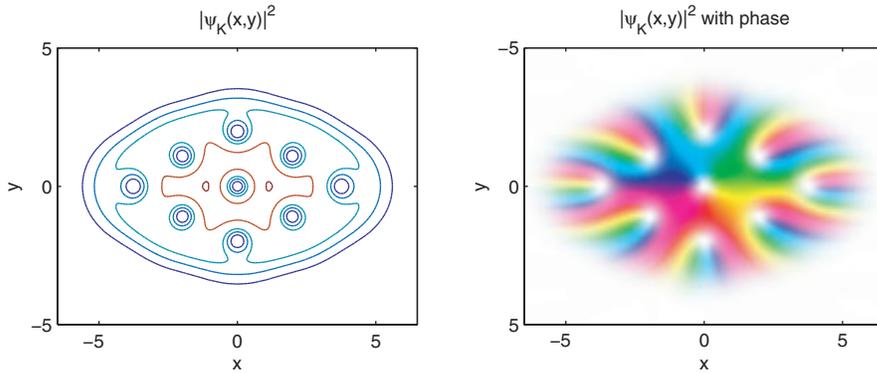


FIG. 6. *Left: visualization of the particle density $|\psi_K(x,y)|^2$ of the computed solution. The contour levels are selected as $0.05 \cdot 10^{-3}$, $0.10 \cdot 10^{-3}$, \dots , $0.30 \cdot 10^{-3}$; the outer contours have smaller values. Right: visualization colored according to density and phase. The color hue is chosen from the standard color wheel based on the phase angle $-i \log(\psi_K/|\psi_K|)$. The corresponding eigenvalue approximation is $\lambda_K = 6.469449$. The solution exhibits several vortices arranged in a regular pattern, as expected for a rotating BEC.*

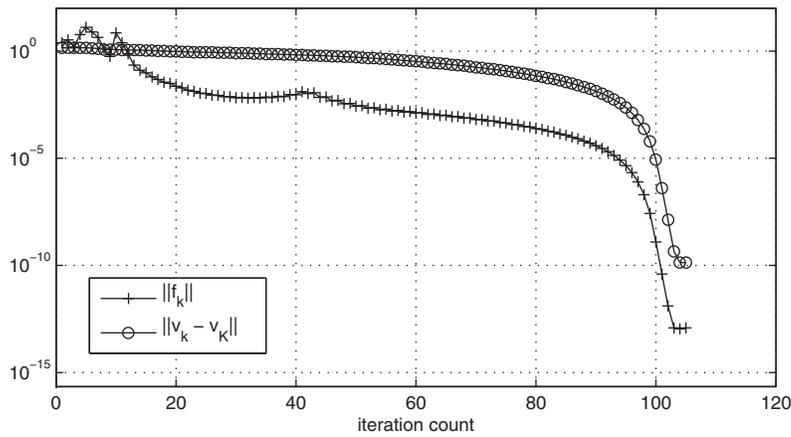
like $\|v_k - v_K\| \approx \|y(t_k) - y_*\| \propto C \exp(-at_k)$ in an asymptotic regime. This can indeed be confirmed in Figure 7. The estimate clearly indicates that $\|v_{k+1} - v_K\|/\|v_k - v_K\|$ should converge when $t_{k+1} = t_k + h_{\max}$. Hence, with the assumption that σ is approximately constant in the regime where we take step length h_{\max} we have $\sigma \approx \lambda_* - 1/h_{\max}$. We can compute the theoretical convergence factor for this σ by using (3.11) and computing μ_2 with the Arnoldi method (and a matrix-vector product from Theorem 5.1). The theoretical convergence factor is visualized together with the estimated convergence factor $\|v_{k+1} - v_K\|/\|v_k - v_K\|$ in Figure 9. The theoretical convergence factor is confirmed for two different choices of h_{\max} .

The heuristic choice of σ is visualized in Figure 9. With the crude estimation of the relation (4.8), $h \approx 1/(\lambda_* - \sigma)$, we see that the step length in the beginning is chosen large, in an intermediate phase it is chosen around the order of magnitude 10, and in the final phase it is chosen larger, and σ is again eventually almost constant.

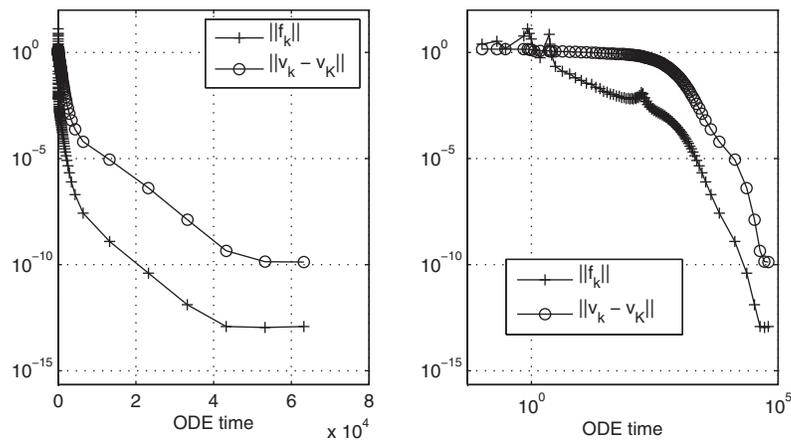
6. Concluding remarks. The favorable convergence properties of the celebrated inverse iteration algorithm for the standard eigenvalue problem are well understood. An important point in this paper is that the generalization we have presented has many of the favorable properties that are present in the inverse iteration algorithm for standard eigenvalue problems. This holds in particular for local convergence and the interpretation as an ODE. We have also illustrated the usefulness of the algorithm by adapting it to a variant of the Schrödinger equation.

The connection between the GPE and the use of inverse iteration presented in this paper has further indirect value. For instance, the tremendous amount of understanding that is available for inverse iteration (for the standard eigenvalue problem) now has the potential to be exploited or adapted to this type of nonlinearity.

Appendix A. Derivation of local error and step length for a variant of Rosenbrock–Euler. The error of the Rosenbrock–Euler method has been studied in the context of the Runge–Kutta methods (e.g., [12, Chapter IV.7]). We need a more specialized result and will for completeness provide an error estimate for the Rosenbrock–Euler method in our setting. Consider the autonomous ODE $y'(t) = f(y(t))$ with $\|y(0)\| = 1$. Suppose the ODE has a structure such that the norm is an



(a)



(b)

FIG. 7. Plot of norm of residual and absolute error as function of iteration count (a) and ODE time t (b). The horizontal axis is linear in the left panel, and logarithmic in the right panel. In the final iterations, the time step is $h_{\max} = 10^4$. The vector f_k denotes the residual, i.e., $f_k = p(v_k)v_k - A(v_k)v_k$.

invariant, i.e., $\|y(t)\| = 1$ for all $t > 0$. Let \tilde{y}_1 be one step of the Rosenbrock–Euler method. By using a Taylor expansion, it is straightforward to show that the local error of the Rosenbrock–Euler step is given by $\tilde{y}_1 - y(h) = h^2q + \frac{h^3}{6}\hat{y}'''(\tau)$, where

$$(A.1) \quad q := -\frac{1}{2}(I - hf'(\tilde{y}_0))^{-1}(I + hf'(\tilde{y}_0))f'(\tilde{y}_0)f(\tilde{y}_0).$$

and $\hat{y}'''(\tau) = (y'''(\tau_1)_1, \dots, y'''(\tau_n)_n)^T$ with $\tau_1, \dots, \tau_n \in [0, h]$. We approximate $\|q\| \approx \frac{1}{2}\|f'(\tilde{y}_0)f(\tilde{y}_0)\|$ and neglect the final term which leads us to the following choice of step length for a given error tolerance:

$$(A.2) \quad h = \sqrt{\frac{2\varepsilon}{\|f'(\tilde{y}_0)f(\tilde{y}_0)\|}}.$$

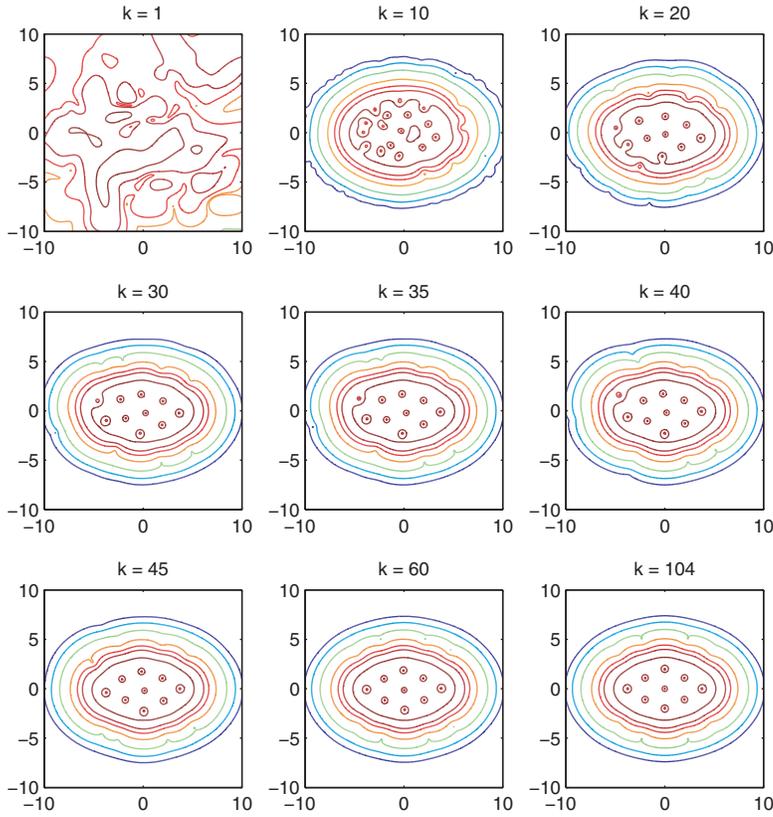


FIG. 8. The approximations $|\psi_k(x, y)|$ for different iterates. The random starting value first turns into the general shape of the solution in ~ 10 iterations, and in the remaining iterations, only the vortices are modified. The contour levels were selected as 10^{-10} , 10^{-8} , 10^{-6} , 10^{-4} , 10^{-3} , $10^{-2.5}$, 10^{-2} , $10^{-1.75}$, $10^{-1.5}$, and $10^{-1.25}$. The outer contours have the smaller values.

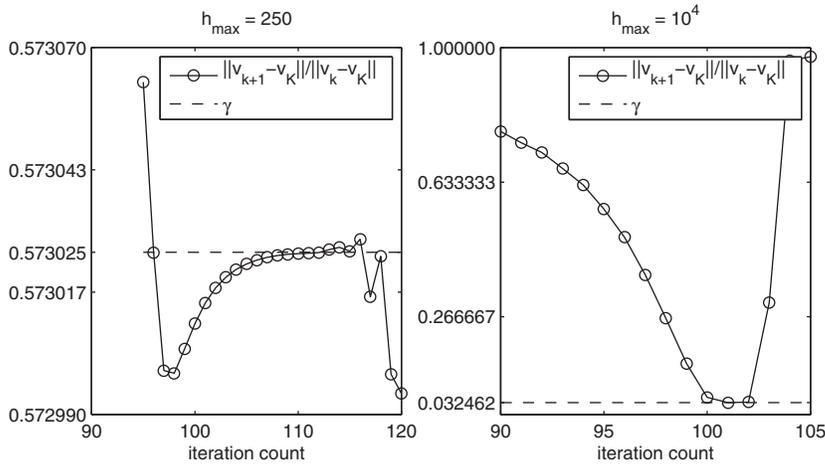


FIG. 9. Estimation of the convergence factor from a calculation with $h_{\max} = 250$ and comparison with the theoretical convergence factor γ (horizontal lines). The asymptotic region starts at iteration $k \approx 95$. The deviations for large k are due to numerical noise in the error estimate $\|v_k - v_K\|$.

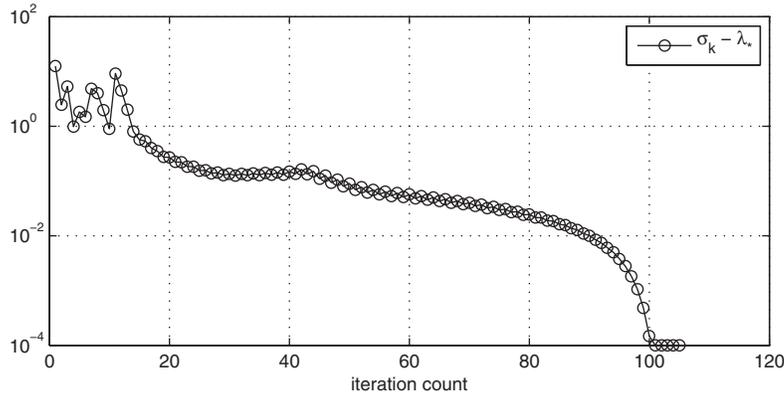


FIG. 10. Plot of $\sigma_k - \lambda_*$, which quickly becomes small and negative.

Acknowledgments. We thank Prof. Tobias Damm (TU Kaiserslautern), Prof. Alexander Ostermann (Univ. Innsbruck), Melina Freitag (Univ. Bath) for comments on preliminary results of this paper. We also thank the referees for providing very constructive comments. In particular, the connection with Newton's methods described in section 3.4 was pointed out by a referee.

REFERENCES

- [1] W. BAO AND Q. DU, *Computing the ground state solution of Bose-Einstein condensates by a normalized gradient flow*, SIAM J. Sci. Comput., 25 (2004), pp. 1674–1697.
- [2] W. BAO AND W. TANG, *Ground-state solution of Bose-Einstein condensate by directly minimizing the energy functional*, J. Comput. Phys., 187 (2003), pp. 230–254.
- [3] A. BLOCH, *Hamiltonian and Gradient Flows, Algorithms and Control*, Fields Inst. Commun. 3, AMS, Providence, RI, 1994.
- [4] M. CALIARI, A. OSTERMANN, S. RAINER, AND M. THALHAMMER, *A minimisation approach for computing the ground state of Gross-Pitaevskii systems*, J. Comput. Phys., 228 (2009), pp. 349–360.
- [5] Y.-S. CHOI, I. KOLTRACHT, P. J. MCKENNA, AND N. SAVYTSKA, *Global monotone convergence of Newton iteration for a nonlinear eigen-problem*, Linear Algebra Appl., 357 (2002), pp. 217–228.
- [6] M. T. CHU, *On the continuous realization of iterative processes*, SIAM Rev., 30 (1988), pp. 375–387.
- [7] Y.-M. J. DEMOULIN AND Y. M. CHEN, *An iteration method for solving nonlinear eigenvalue problems*, SIAM J. Appl. Math., 28 (1975), pp. 588–595.
- [8] A. EDELMAN, T. A. ARIAS, AND S. T. SMITH, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Anal. Appl., 20 (1998), pp. 303–353.
- [9] M. FREITAG AND A. SPENCE, *Convergence of inexact inverse iteration with application to preconditioned iterative solves*, BIT, 47 (2007), pp. 27–44.
- [10] G. GOLUB AND C. VAN LOAN, *Matrix Computations*, Johns Hopkins University Press, Baltimore, MD, 2007.
- [11] E. HAIRER, C. LUBICH, AND G. WANNER, *Geometric Numerical Integration. Structure-Preserving Algorithms for Ordinary Differential Equations*, 2nd ed., Springer Ser. Comput. Math. 31, Springer, Berlin, 2006.
- [12] E. HAIRER AND G. WANNER, *Solving Ordinary Differential Equations. II: Stiff and Differential-Algebraic Problems*, 2nd rev. ed., Springer Ser. Comput. Math. 14, Springer, Berlin, 1996.
- [13] D. HINRICHSSEN AND A. J. PRITCHARD, *Mathematical Systems Theory I*, Springer, Berlin, 2005.
- [14] I. C. F. IPSEN, *Computing an eigenvector with inverse iteration*, SIAM Rev., 39 (1997), pp. 254–291.
- [15] E. JARLEBRING, *Convergence factors of Newton methods for nonlinear eigenvalue problems*, Linear Algebra Appl., 436 (2012), pp. 3943–3953.

- [16] E. JARLEBRING AND W. MICHIELS, *Analyzing the convergence factor of residual inverse iteration*, BIT, 51 (2011), pp. 937–957.
- [17] A. V. KNYAZEV AND K. NEYMEYR, *Gradient flow approach to geometric convergence analysis of preconditioned eigensolvers*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 621–628.
- [18] L. LANDAU AND E. LIFSHITZ, *Quantum Mechanics: Non-relativistic Theory*, Pergamon Press, Oxford, 1965.
- [19] R. MAHONY AND P.-A. ABSIL, *The continuous-time Rayleigh quotient flow on the sphere*, Linear Algebra Appl., 368 (2003), pp. 343–357.
- [20] V. MEHRMANN AND H. VOSS, *Nonlinear eigenvalue problems: A challenge for modern eigenvalue methods*, GAMM Mitt., Ges. Angew. Math. Mech., 27 (2004), pp. 121–152.
- [21] R. MEYER, *Nonlinear eigenvector algorithms for local optimization in multivariate data analysis*, Linear Algebra Appl., 264 (1997), pp. 225–246.
- [22] A. NEUMAIER, *Residual inverse iteration for the nonlinear eigenvalue problem*, SIAM J. Numer. Anal., 22 (1985), pp. 914–923.
- [23] K. NEYMEYR, *A note on inverse iteration*, Numer. Linear Algebra Appl., 12 (2005), pp. 1–8.
- [24] K. NEYMEYR, *A geometric theory for preconditioned inverse iteration IV: On the fastest convergence cases*, Linear Algebra Appl., 415 (2006), pp. 114–139.
- [25] J. ORTEGA AND W. RHEINBOLDT, *Iterative Solution of Nonlinear Equations in Several Variables*, Classics Appl. Math. 30, SIAM, Philadelphia, 2000.
- [26] A. OSTROWSKI, *On the convergence of the Rayleigh quotient iteration for the computation of the characteristic roots and vectors. I*, Arch. Ration. Mech. Anal., 1 (1959), pp. 233–241.
- [27] G. PETERS AND J. H. WILKINSON, *Inverse iteration, ill-conditioned equations and Newton's method*, SIAM Rev., 21 (1979), pp. 339–360.
- [28] V. SIMONCINI AND L. ELDÉN, *Inexact Rayleigh quotient-type methods for eigenvalue computations*, BIT, 42 (2002), pp. 159–182.
- [29] H. VOSS, *Nonlinear eigenvalue Problems*, in Handbook of Linear Algebra, 2nd ed., J. Hogben, ed., CRC Press Ser. Discrete Math. Appl. 164, Boca Raton, FL 2013.
- [30] D. S. WATKINS, *Isospectral flows*, SIAM Rev., 26 (1984), pp. 379–391.
- [31] D. S. WATKINS AND L. ELSNER, *Self-similar flows*, Linear Algebra Appl., 110 (1988), pp. 213–242.
- [32] W.-Y. YAN, U. HELMKE, AND J. B. MOORE, *A global analysis of Oja's flow for neural networks*, IEEE Trans. Neural Networks, 1994.