

Solving the Discrete Euler-Arnold Equations for the Generalized Rigid Body Motion

João R. Cardoso*

*Coimbra Polytechnic–ISEC, and
Center for Mathematics, University of Coimbra, Portugal.*

Pedro Miraldo

*Institute for Systems and Robotics (LARSyS),
Instituto Superior Técnico, University of Lisbon, Portugal.*

Abstract

We propose three iterative methods for solving the Moser-Veselov equation, which arises in the discretization of the Euler-Arnold differential equations governing the motion of a generalized rigid body. We start by formulating the problem as an optimization problem with orthogonal constraints and proving that the objective function is convex. Then, using techniques from optimization on Riemannian manifolds, the three feasible algorithms are designed. The first one splits the orthogonal constraints using the Bregman method, whereas the other two methods are of the steepest-descent type. The second method uses the Cayley-transform to preserve the constraints and a Barzilai-Borwein step size, while the third one involves geodesics, with the step size computed by Armijo’s rule. Finally, a set of numerical experiments are carried out to compare the performance of the proposed algorithms, suggesting that the first algorithm has the best performance in terms of accuracy and number of iterations. An essential advantage of these iterative methods is that they work even when the conditions for applicability of the direct methods available in the literature are not satisfied.

Keywords: Discrete Euler-Arnold equations, matrix equation, Moser-Veselov equation, optimization with orthogonal constraints, orthogonal matrices, skew-symmetric matrices.

*Corresponding author: João R. Cardoso

Email addresses: `jocar@isec.pt` (João R. Cardoso), `pedro.miraldo@tecnico.ulisboa.pt` (Pedro Miraldo)

1. INTRODUCTION

In [29], Moser and Veselov proposed the following equations to discretize the classical Euler-Arnold differential equations for the motion of a generalized rigid body:

$$\begin{aligned} M_{k+1} &= \omega_k M_k \omega_k^T \\ M_k &= \omega_k^T J - J \omega_k, \end{aligned} \tag{1}$$

where M_k is the angular momentum with respect to the body (here represented by a skew-symmetric matrix), J is the inertia matrix (symmetric positive definite), and ω_k (orthogonal matrix) is the angular velocity. Rigid body equations arise in several applications, e.g., celestial mechanics, molecular dynamics, mechanical robotics, and flight control, where they are used in particular to understand the body-body interactions of particles like planets, atoms, and molecules. See, for instance, [5, 13, 23, 31, 20, 21] and the references therein.

The main challenge of solving (1) is to find an orthogonal matrix ω_k in the second equation, by assuming that J and M_k are given. Mathematically, the problem consists of finding an orthogonal matrix X (for convenience, here we use X instead of ω) such that

$$XJ - JX^T = M, \tag{2}$$

where J is a given symmetric positive definite matrix, and M is a known skew-symmetric matrix. All the matrices involved are square of order n . The matrix equation (2) is known as the *Moser-Veselov* equation and was firstly investigated in [29], where the authors based their developments on factorizations of certain matrix polynomials. A different approach, but computationally more efficient, was provided later in [8], where the authors noted that (2) can be connected with a certain algebraic Riccati equation and, in turn, with the *Hamiltonian* matrix

$$\mathcal{H} = \begin{bmatrix} M/2 & I \\ M^2/4 + J^2 & M/2 \end{bmatrix}. \tag{3}$$

We now revisit some results stated in [8], concerning the existence and uniqueness of solutions of (2).

Theorem 1. *For the matrix equation (2):*

1. *There exists a solution $X \in \mathcal{SO}(n)$ (the special orthogonal or rotation group of order n) if and only if the size of the Jordan blocks associated to the pure imaginary eigenvalues of \mathcal{H} (if any) is even;*
2. *(2) has a unique solution $X \in \mathcal{SO}(n)$ if and only if the spectrum of \mathcal{H} is pure imaginary and the size of the Jordan blocks associated to each (nonzero) eigenvalue is even.*

From Theorem 1, we can see that the Moser-Veselov equation may have no solution in $\mathcal{SO}(n)$ if the associated Hamiltonian matrix \mathcal{H} has any pure imaginary eigenvalue with a Jordan block of odd size. It is also known that the existence of purely imaginary eigenvalues in \mathcal{H} causes significant difficulties in solving (2). To avoid those situations, as far as we know, in all the existing algorithms for solving the equation, it is assumed a priori that \mathcal{H} does not admit any pure imaginary eigenvalue (see, [27, Sec. 1.2], [29, Sec. 1.4], and [31, Sec. G]). Moreover, the algorithms based on solving the associated algebraic Riccati equation require the strong condition that the matrix $M^2/4 + J^2$ must be symmetric positive definite. These issues have motivated us to investigate methods whose applicability does not require those restrictive conditions. As we can see later in Sec. 5, the three proposed optimization algorithms produce special orthogonal solutions, even when $M^2/4 + J^2$ is not symmetric positive definite. Those iterative algorithms may also be used in problems where \mathcal{H} has purely imaginary eigenvalues associated with Jordan blocks of even size (check Theorem 1) but, as will be illustrated later in Sec. 5, the convergence may slow down.

Problem 2. *Let $\|\cdot\|_F$ denote the Frobenius norm, i.e., $\|A\|_F := \sqrt{\text{trace}(A^T A)}$. The problem of finding a special orthogonal solution X in (2) can be formulated as an optimization problem in the following way:*

$$\min_{X \in \mathcal{SO}(n)} \|XJ - JX^T - M\|_F^2. \quad (4)$$

In Problem 2, we have chosen the Frobenius norm because its definition in terms of the trace of a matrix allows us to access the derivatives of the objective function easily, making it more suitable to handle optimization problems than other norms, like, for instance, spectral or infinity norms.

The literature on numerical methods for solving non-linear constrained problems, like (4), is large; see for instance [30, 24, 33, 28, 7]. However, due to the complicated expression of the

objective function and the large number of constraints arising from the conditions $X^T X = I$ and $\det(X) = 1$, some care must be taken with the choice of the methods.

Techniques from Riemannian geometry for solving optimization problems with orthogonal constraints have attracted the interest of many researchers in the last decades; see [2, 14], and the references therein. An essential feature of those techniques is that they allow the transformation of a constrained optimization problem into an unconstrained one. Moreover, since the set of orthogonal matrices is a manifold and provided that the objective function satisfies some smoothness requirements, we can make available tools such as Euclidean gradients, Riemannian gradients, retractions, and geodesics.

The three methods presented in this work evolve on the orthogonal manifold and belong to the family of line search methods on manifolds described in [2, Ch. 4]. They are iterative and feasible (or constraint-preserving), in the sense that, starting with a matrix $X_0 \in \mathcal{SO}(n)$, all the iterates X_k also stay in $\mathcal{SO}(n)$.

The major contributions of this work are:

- the development of three effective algorithms for solving Problem 2 and in turn the matrix equation (2) (i.e., the Moser-Veselov equation), that do not require the condition $M^2/4 + J^2 > 0$;
- a detailed discussion on the properties of the optimization problems, including convexity issues;
- a novel approach for solving the unconstrained optimization problems arising in each iteration of the Bregman splitting algorithm and a discussion about the reasons that led the MATLAB's `fminunc` function to give unsatisfactory results in some circumstances;
- a novel relative residual capable to infer about the quality of the computed solution of the Moser-Veselov equation;
- careful modifications on existing algorithms for solving optimization problems with orthogonal constraints, to make them suitable for our particular problems.

In the next section, we derived a workable expression to the objective function of Problem 2. Sec. 3 presents the three algorithms proposed in this paper. In Sec. 4, numerical issues of the

algorithms are discussed, and, in Sec. 5, a selection of experiments are carried out to illustrate the performance of the algorithms. In Sec. 6, some conclusions will be drawn.

2. Rewriting the Objective Function

Let us denote by $F(X) := \|XJ - JX^T - M\|_F^2$ the objective function arising in (4). Using the properties of the trace of a matrix and attending that $J^T = J$ and $M^T = -M$, we have (detailed calculation is omitted):

$$\begin{aligned} F(X) &= \text{trace}((XJ - JX^T - M)^T(XJ - JX^T - M)) \\ &= 2 \text{trace}(JX^T XJ) - 2 \text{trace}(XJXJ) + 4 \text{trace}(MXJ) - \text{trace}(M^2). \end{aligned} \quad (5)$$

Now, if we take into account the orthogonality of X , that is, $X^T X = X X^T = I$, then $F(X)$ can be simplified to (a different notation is used):

$$\tilde{F}(X) = -2 \text{trace}((JX)^2) + 4 \text{trace}(XJM) + 2 \text{trace}(J^2) - \text{trace}(M^2), \quad (6)$$

which is the restriction of $F(X)$ to the orthogonal group $\mathcal{O}(n)$, that is, $F(X) = \tilde{F}(X)$, for any $X \in \mathcal{O}(n)$, but, in general, $F(X) \neq \tilde{F}(X)$, if $X \notin \mathcal{O}(n)$. Hence, the problem (4) may be simplified to:

$$\min_{X \in \mathcal{SO}(n)} \tilde{F}(X) = -2 \text{trace}((JX)^2) + 4 \text{trace}(XJM) + \alpha, \quad (7)$$

where $\alpha := 2 \text{trace}(J^2) - \text{trace}(M^2)$.

If x_{ij} denotes the entry (i, j) of the matrix X , then both $F(X)$ defined in (5) and $\tilde{F}(X)$ in (6) are differentiable functions in \mathbb{R}^{n^2} , because they are quadratic polynomials in the n^2 variables x_{ij} . In the following lemma we show that $F(X)$ is convex in the set of all $n \times n$ matrices with real entries $\mathbb{R}^{n \times n}$, that is:

$$F(tX_1 + (1-t)X_2) \leq tF(X_1) + (1-t)F(X_2), \quad (8)$$

for any $t \in [0, 1]$ and $X_1, X_2 \in \mathbb{R}^{n \times n}$.

Lemma 3. *The function $F(X)$ given in (5) is convex in $\mathbb{R}^{n \times n}$.*

Proof. Let us denote $f(X) := \|XJ - JX^T - M\|_F$. Attending that the Frobenius norm satisfies the triangle inequality, we have

$$\begin{aligned}
f(tX_1 + (1-t)X_2) &= \left\| (tX_1 + (1-t)X_2)J - J(tX_1 + (1-t)X_2)^T - M \right\|_F \\
&= \left\| t(X_1J - JX_1^T) + (1-t)(X_2J - JX_2^T) - M \right\|_F \\
&= \left\| t(X_1J - JX_1^T) + (1-t)(X_2J - JX_2^T) - (t+1-t)M \right\|_F \\
&= \left\| t(X_1J - JX_1^T - M) + (1-t)(X_2J - JX_2^T - M) \right\|_F \\
&\leq t \|X_1J - JX_1^T - M\|_F + (1-t) \|X_2J - JX_2^T - M\|_F \\
&= tf(X_1) + (1-t)f(X_2), \tag{9}
\end{aligned}$$

for all $t \in [0, 1]$ and $X_1, X_2 \in \mathbb{R}^{n \times n}$. Consider the scalar function $g(y) = y^2$. Since $f(X) \geq 0$, for any $X \in \mathbb{R}^{n \times n}$, and g is non-decreasing in $[0, +\infty[$, we conclude that $F(X) = g(f(X))$ is convex. \square

Similarly, we could show that $\tilde{F}(X)$ is convex in $\mathbb{R}^{n \times n}$. Note, however, that the constraints of the optimization problem (7) are non-convex, that is, for $P, Q \in \mathcal{SO}(n)$ and $t \in [0, 1]$, in general $tP + (1-t)Q \notin \mathcal{SO}(n)$, which makes the problem much more difficult.

Now, we use the rules for the derivatives of the trace function (see, for instance, [25, Ch. 10]) to obtain the expressions of the Euclidean gradients (derivatives with respect to X) of those functions:

$$\nabla F(X) = 4XJ^2 - 4JX^TJ - 4MJ \tag{10}$$

$$\nabla \tilde{F}(X) = -4JX^TX - 4MJ. \tag{11}$$

The Riemannian gradients in the orthogonal manifold can be defined by:

$$\text{grad } F(X) = \nabla F(X) X^T - X \nabla F(X)^T \tag{12}$$

$$\text{grad } \tilde{F}(X) = \nabla \tilde{F}(X) X^T - X \nabla \tilde{F}(X)^T. \tag{13}$$

We recall that these Riemannian gradients belong to the orthogonal group's tangent space, which is the set of skew-symmetric matrices (see, for instance, [14]).

In the next section, we propose three algorithms to solve (2).

3. Algorithms

There is a vast literature on methods for optimizing functions on the orthogonal group or, more generally, on the Stiefel manifold (e.g., [1, 14, 15, 19, 22, 26, 35, 36]). Among those methods, we have selected three state-of-the-art ones that we believe to be well suited for our specific objective function. The first algorithm (Section 3.1) splits the orthogonal constraints in a Bregman’s style [6, 32] and is based on the SOC algorithm proposed by Lai and Osher in [22, Alg. 2]. Our second algorithm (Section 3.2) is inspired in the feasible method developed by Wen and Yin in [35, Alg. 2], which uses a retraction — the Cayley transform — instead of geodesics. It is a line search method and, to find the appropriate step size, the method of Barzilai-Borwein (BB) [4] is used. The third algorithm involves line search techniques, namely the Armijo’s rule, and is based on a proposal by Abrudan *et al.* in [1, Table II]. It is of steepest descent type and involves geodesics, more specifically matrix exponentials.

3.1. Algorithm Based on Bregman Splitting

The SOC Algorithm [22, Alg. 2] applied to Problem 2 is summarized in the following steps:

1. Choose a positive scalar r and a starting matrix X_0 . Set $P_0 = X_0$ and $B_0 = 0$;
2. While “not converge” do
 - (a) $X_k = \underset{X}{\operatorname{argmin}} F(X) + \frac{r}{2} \|X - P_{k-1} + B_{k-1}\|_F^2$, where $F(X)$ is defined by (5);
 - (b) $Y_k \leftarrow X_k + B_{k-1}$;
 - (c) Compute the singular value factorization: $Y_k = UDV^T$;
 - (d) $P_k \leftarrow UV^T$;
 - (e) $B_k \leftarrow B_{k-1} + X_k - P_k$.

A drawback of this algorithm is the requirement of solving the (unconstrained) optimization problem in Step 2(a). Let us write down the corresponding objective function in terms of the trace. Set $G_k(X) := \frac{r}{2} \|X - P_{k-1} + B_{k-1}\|_F^2$ and $C_k = -P_{k-1} + B_{k-1}$. After a few calculations, we have

$$G_k(X) = \frac{r}{2} (\operatorname{trace}(X^T X) + 2\operatorname{trace}(C_k^T X) + \operatorname{trace}(C_k^T C_k)). \quad (14)$$

Hence, the objective function in Step 2(a) is $\mathcal{F}_k(X) := F(X) + G_k(X)$, where $F(X)$ denotes the function given in (5), and the associated unconstrained optimization problem may be

formulated as

$$\min_X \mathcal{F}_k(X) = \min_X 2 \operatorname{trace}(JX^T XJ) - 2 \operatorname{trace}(XJXJ) + 4 \operatorname{trace}(MXJ) - \operatorname{trace}(M^2) + \frac{r}{2} (\operatorname{trace}(X^T X) + 2 \operatorname{trace}(C_k^T X) + \operatorname{trace}(C_k^T C_k)), \quad (15)$$

where C_k, J, M are given square matrices of order n , J and M are, respectively, symmetric positive definite and skew-symmetric, and r is a positive parameter. By a similar argument to the one used in Lemma 3, we can show that, for each k , $G_k(X)$ in (14) is convex, and the same is valid to the objective function $\mathcal{F}_k(X)$. This is a very useful property because it guarantees that local minima are global as well.

Since each iteration k of the above SOC algorithm requires the solution of a convex unconstrained optimization problem of the form (15), whose objective function changes for each k , according to the entries of the matrix C_k , a possible approach to solve each one is to use the MATLAB's function `fminunc`, which is based on quasi-Newton and trust region methods. More precisely, `fminunc` is based, by default, on the Broyden, Fletcher, Goldfarb, and Shanno quasi-Newton method, which is also known as the BFGS method (check [30, Sec. 8.1] and the references therein). However, if the gradient $\nabla \mathcal{F}(X_k)$ is provided, then `fminunc` switches to a trust region method based on the proposals of Coleman and Li [10, 11].

We have solved many unconstrained problems of the form (15) using `fminunc` but, despite the convexity, the results were not so good as expected, either in terms of speed of convergence or in terms of accuracy. We learnt from our experiments that `fminunc` (without gradient and with the default tolerance of 10^{-6}) is reliable only for very small size problems (say, $n \leq 5$). As n increases, we observed that in some iterations k of the SOC algorithm, `fminunc` was unable to minimize $\mathcal{F}(X_k)$. It displayed warnings like "local minimum possible" or "solver stopped prematurely". Recall that finding the minimum of (15) in all the iterations is necessary to guarantee the convergence of the SOC algorithm. For a given fixed k , let $X_k^{(i)}$ denotes the i -th iteration generated by `fminunc` when applied to the minimization $\mathcal{F}(X_k)$. An interesting fact we have observed was that, as i increased, $\mathcal{F}(X_k^{(i)})$ decreased quite fast to values lower than 10^{-6} , while the components of the gradient vector $\nabla \mathcal{F}(X_k^{(i)})$ decreased in a slow fashion towards zero. This implies that `fminunc` involves a large number of iterations to guarantee that the norm of the gradient is lower than a fixed tolerance and hence that $X_k^{(i)}$ satisfies (up

to that tolerance) the first-order necessary conditions. Recall that, if X_* is a local (or global) minimizer of and \mathcal{F} (which is continuously differentiable), we must have $\nabla\mathcal{F}(X_*) = 0$.

However, if the gradient function $\nabla\mathcal{F}$ (check (16)) is provided to `fminunc` and the tolerance is set to 10^{-5} , we see that its performance improves and, as shown later in Section 5, for equations of small size (say, $n \leq 15$), the usage of `fminunc` can be viewed as a possible approach to make the SOC algorithm effective.

To deal with smaller tolerances and equations involving matrices with larger size, we propose a different approach, which is described next. As we will see later, this approach seems to be very promising, either in terms of accuracy and computational cost.

We start by finding the zeros of the gradient of $\mathcal{F}(X)$ (here, to simplify the notation, we omit the subscript k). We note that the expression of $\mathcal{F}(X)$ is non-linear and involves n^2 variables.

A few calculations lead to the following expression to the gradient of the objective function $\mathcal{F}(X)$:

$$\nabla\mathcal{F}(X) = 4XJ^2 - 4JX^TJ - 4MJ + r(X + C). \quad (16)$$

We know that local minima of $\mathcal{F}(X)$ are among the zeros of its Euclidean gradient. Since \mathcal{F} is convex, those local minima (if any) will be global as well. Therefore, we need to investigate the solutions of the matrix equation $\nabla\mathcal{F}(X) = 0$, which is equivalent to

$$X(4J^2 + rI) - 4JX^TJ = 4MJ - rC. \quad (17)$$

An easy way of solving (17) is achieved by performing vectorization. Let $\text{vec}(\cdot)$ stand for the operator that stacks the columns of a matrix $n \times n$ into a long vector of size $n^2 \times 1$, and let \otimes denote the Kronecker product. It is well-known that $\text{vec}(AYB) = (B^T \otimes A) \text{vec}(Y)$ and that

$$\text{vec}(A^T) = \Pi \text{vec}(A), \quad (18)$$

where Π is the commutation (or permutation) matrix of order $n^2 \times n^2$ (check [25, Ch. 7, Sec.

9.2]). Applying the vec operator to (17) yields

$$\begin{aligned}
& \text{vec}(X(4J^2 + rI)) - 4 \text{vec}(JX^T J) = \text{vec}(4MJ - rC) \iff \\
& ((4J^2 + rI) \otimes I) \text{vec}(X) - 4(J \otimes J) \text{vec}(X^T) = \text{vec}(4MJ - rC) \iff \\
& ((4J^2 + rI) \otimes I) \text{vec}(X) - 4(J \otimes J)\Pi \text{vec}(X) = \text{vec}(4MJ - rC) \iff \\
& [(4J^2 + rI) \otimes I - 4(J \otimes J)\Pi] \text{vec}(X) = \text{vec}(4MJ - rC). \tag{19}
\end{aligned}$$

Note that (19) corresponds to a linear system of type

$$A\mathbf{x} = \mathbf{b},$$

where $A := (4J^2 + rI) \otimes I - 4(J \otimes J)\Pi$ is an $n^2 \times n^2$ matrix, $\mathbf{x} := \text{vec}(X)$, and $\mathbf{b} := \text{vec}(4MJ - rC)$. The vectorization approach is very useful to understand the theory of the matrix equation (17). If A is non-singular, a unique solution to (17) is guaranteed. Otherwise, if A is singular, (17) may have infinitely many solutions or no solutions. In the numerical examples we have considered (including the ones to be shown in Section 5), the matrix A encountered was always non-singular.

It turns out, however, that solving (17) through the linear system $A\mathbf{x} = \mathbf{b}$ would require $O(n^6)$ operations, which is prohibitive, especially if n large. This has motivated us to look for less expensive methods for solving (17).

Since J is assumed to be non-singular, a right multiplication of the matrix equation (17) by J^{-1} yields

$$X(4J + rJ^{-1}) - 4JX^T = 4M - rCJ^{-1}. \tag{20}$$

Setting $Y := X^T$, $A_1 := -4J$, $A_2 := 4J + rJ^{-1}$, and $A_3 := 4M - rCJ^{-1}$, (20) can be rewritten in the form

$$A_1 Y + Y^T A_2 = A_3, \tag{21}$$

which is a Sylvester-type equation. An effective method to solve (21) may be found in [34, Alg. 3.1], which is based on the so-called QZ decomposition. Do not confuse (21) with the more easily to handle classical Sylvester matrix equation $A_1 X + X A_2 = A_3$. Although the QZ decomposition is quite expensive, it can be performed in $O(n^3)$ operations. Hence, it is possible to solve (17) and the Step 2(a) of the SOC algorithm efficiently by $O(n^3)$ operations, which is the typical cost for problems involving matrix-matrix products.

Algorithm 3.1 Bregman splitting algorithm to solve the Moser-Veselov equation (2).

- 1: Choose a positive scalar r and a starting matrix $X_0 \in \mathcal{SO}(n)$;
 - 2: $P_0 \leftarrow X_0$; $B_0 \leftarrow 0$;
 - 3: $J_1 \leftarrow J_1^{-1}$; $A_1 \leftarrow -4J$; $A_2 \leftarrow 4J + rJ_1$;
 - 4: $k \leftarrow 1$;
 - 5: **while** “not converge” **do**
 - 6: $A_3 \leftarrow 4M - r(B_{k-1} - P_{k-1})J_1$;
 - 7: Solve $A_1X_k + X_k^T A_2 = A_3$ for X_k , using [34, Algs. 3.1];
 - 8: $X_k \leftarrow X_k^T$;
 - 9: $Y_k \leftarrow X_k + B_{k-1}$;
 - 10: Compute the SVD $Y_k = UDV^T$ and set $P_k \leftarrow UV^T$;
 - 11: Compute the SVD $X_k = UDV^T$ and set $X_k \leftarrow UV^T$;
 - 12: $B_k \leftarrow B_{k-1} + X_k - P_k$;
 - 13: $k \leftarrow k + 1$.
-

A drawback we have noticed when performing experiments with a direct application of the algorithm of [22], to solve the Moser-Veselov equation, was the poor feasibility of the approximation obtained when compared to the other algorithms to be addressed in the next sections. That is, denoting by \tilde{X} the approximation obtained for the solution, we have observed that the value $\|\tilde{X}^T \tilde{X} - I\|_F$ was not close enough to zero. To overcome this issue, we project X_k onto the orthogonal group by computing its singular value decomposition (SVD). Provided that (21) has a unique solution Y_* , that means that $X_* = Y_*^T$ is the unique minimizer of the objective function in (15).

Algorithm 3.1 summarizes the main steps of the Bregman splitting algorithm to solve the Moser-Veselov equation, (2). The next subsection presents two alternative approaches for solving this problem.

3.2. Two Steepest Descent-Type Algorithms

A successful method to solve optimization problems with orthogonal constraints is the feasible iterative method developed in [35, Alg. 2]. At each iteration, the skew-symmetric Riemannian gradient (13) is multiplied by a suitable positive number τ (step-size), and transformed

into an orthogonal matrix by means of the Cayley transformation

$$Y(\tau) = \left(I + \frac{\tau}{2}W\right)^{-1} \left(I - \frac{\tau}{2}W\right) X, \quad (22)$$

where W is the Riemannian gradient and X is an orthogonal matrix. In the iterative procedure, the orthogonal matrix $Y(\tau)$ may be viewed as an improvement of a previous approximation X .

There are many methods available to compute the step-size τ . In [35, Alg. 2], the authors recommend a non-monotone linear search method, because of its good theoretical properties regarding the convergence. However, it is not considered here because it has led to poor results in many experiments carried out (not shown here) with particular Moser-Veselov equations. In our modified version of the Wen-Yin algorithm [35, Alg. 2], which is presented in Algorithm 3.2, we use, instead, the alternating BB method of [12]:

$$\tau_k = \begin{cases} \|S_{k-1}\|_F^2 / |\langle S_{k-1}, Y_{k-1} \rangle| & \text{if } k \text{ is odd} \\ |\langle S_{k-1}, Y_{k-1} \rangle| / \|Y_{k-1}\|_F^2 & \text{if } k \text{ is even} \end{cases}, \quad (23)$$

where $S_{k-1} := X_k - X_{k-1}$, $Y_{k-1} := \text{grad } F(X_k) - \text{grad } F(X_{k-1})$ and $\langle A, B \rangle = \text{trace}(A^T B)$ denotes the Euclidean scalar product.

The second steepest descent-type algorithm addressed here (Algorithm 3.3) is a variation of the approaches proposed in [26, 1] and we refer the reader to those papers for more technical details. In a few words, Algorithm 3.3 starts with an initial approximation $X_0 \in \mathcal{SO}(n)$, finds the skew-symmetric matrix $\tilde{F}(X_k)$ (the gradient direction on the manifold), and performs several steps along geodesics until convergence. We recall that geodesics on $\mathcal{SO}(n)$ (i.e., curves giving the shortest path between two points in the manifold) can be defined through the matrix exponential as:

$$G(t) = G(0) e^{\mu S},$$

where $S \in \mathbb{R}^{n \times n}$ is a skew-symmetric matrix and μ is a real scalar. In Algorithm 3.3, the positive scalar μ_k controls the length of the ‘‘tangent vector’’ and, in turn, the algorithm’s overall convergence. To find an almost optimal μ_k , the algorithm uses the Armijo’s step-size rule [33, Sec.1.3].

Algorithm 3.2 Algorithm to solve the Moser-Veselov equation (2) inspired on the steepest descent-type method of Wen&Yin [35]. The matrix functions $\tilde{F}(X)$ and $\nabla\tilde{F}(X)$ are defined, respectively, in (6) and (11).

- 1: Choose $\tau > 0$ and a starting matrix $X_0 \in \mathcal{SO}(n)$;
 - 2: $f_0 \leftarrow \tilde{F}(X_0)$; $G_0 \leftarrow \nabla\tilde{F}(X_0)$;
 - 3: $W_0 \leftarrow G_0 X_0^T - X_0 G_0^T$;
 - 4: $k \leftarrow 1$;
 - 5: **while** “not converge” **do**
 - 6: $Y_k \leftarrow (I + 0.5\tau W_{k-1})^{-1}(I - 0.5\tau W_{k-1})X_{k-1}$;
 - 7: $X_k \leftarrow Y_k$;
 - 8: $f_k \leftarrow \tilde{F}(X_k)$ and $G_k \leftarrow \nabla\tilde{F}(X_k)$;
 - 9: $W_k \leftarrow G_k X_k^T - X_k G_k^T$;
 - 10: $S_k \leftarrow X_k - X_{k-1}$;
 - 11: $N_k \leftarrow W_k - W_{k-1}$;
 - 12: **if** k is even **then**
 - 13: $\tau \leftarrow \text{trace}(S_k^T S_k) / |\text{trace}(S_k^T N_k)|$
 - 14: **else**
 - 15: $\tau \leftarrow |\text{trace}(S_k^T N_k)| / \text{trace}(N_k^T N_k)$.
-

4. Numerical Issues

This section addresses some numerical issues associated with the implementation of the three algorithms proposed so far to solve the Moser-Veselov equation.

4.1. Convergence

Assuming that all the pure imaginary eigenvalues (if any) of the $2n \times 2n$ matrix (3) have Jordan blocks with even size, the existence of at least a solution in $\mathcal{SO}(n)$ of the Moser-Veselov equation is guaranteed (check Theorem 1). Provided that a careful choice of the starting matrix X_0 is made, one of those solutions may be obtained by the three proposed algorithms. We recall that finding an initial guess X_0 that minimizes the number of iterations in iterative methods for solving equations is, in general, a challenging problem. However, steepest-descent

Algorithm 3.3 Algorithm to solve the Moser-Veselov equation (2) inspired on the steepest descent-type methods of Manton [26] and Abrudan [1]. The matrix functions $\tilde{F}(X)$ and $\nabla\tilde{F}(X)$ are defined, respectively, in (6) and (11).

```

1:  $X_0 \in \mathcal{SO}(n)$  is an initial guess;
2:  $\mu_1 \leftarrow 1$ ;
3:  $\delta \leftarrow 1$ ;
4:  $\tau \leftarrow \text{tol}$ ;
5:  $k \leftarrow 0$ ;
6: while  $\delta > \tau$  do
7:    $Z_k \leftarrow \nabla\tilde{F}(X_k) X_k^T - X_k \nabla\tilde{F}(X_k)^T$ ;
8:    $z_k \leftarrow 0.5 \text{ trace}(Z_k Z_k^T)$ ;
9:    $P_k \leftarrow \text{expm}(-\mu_k Z_k)$ ;
10:   $Q_k \leftarrow P_k P_k$ ;
11:  while  $\tilde{F}(X_k) - \tilde{F}(Q_k X_k) \geq \mu_k z_k$  do
12:     $P_k \leftarrow Q_k$ ;
13:     $Q_k \leftarrow P_k P_k$ ;
14:     $\mu_k \leftarrow 2\mu_k$ ;
15:  while  $\tilde{F}(X_k) - \tilde{F}(Q_k X_k) < 0.5\mu_k z_k$  do
16:     $P_k \leftarrow \text{expm}(-\mu_k Z_k)$ ;
17:     $\mu_k \leftarrow 0.5\mu_k$ ;
18:   $X_{k+1} \leftarrow P_k X_k$ ;
19:   $\delta \leftarrow \|X_{k+1} - X_k\|_F$ ;
20:   $k \leftarrow k + 1$ ;
21:  $X \leftarrow X_k$ .

```

algorithms combined with suitable methods for computing the step size have good convergence properties (see [4] for the BB method and [33, Sec. 1.3.2] for the Armijo's method). They have linear convergence, but, in contrast to Newton-type methods, in general, it is easier to find an X_0 for which the iterative sequence generated by the method converges.

In our specific case, we have observed through many experiments (some of them will be

shown in Sec. 5) that $X_0 = I$ is a reasonable choice for the three proposed algorithms, in the sense that it leads to convergence towards a special orthogonal solution. Experiments where X_0 has been taken as a randomized special orthogonal matrix will be considered in Sec. 5, but with a slower convergence; see Figures 3 and 4.

4.2. Computational Cost

The three proposed algorithms require $O(n^3)$ operations, which, as written before, is acceptable for algorithms involving matrix computations. However, this information is vague, and we shall give a more sharp estimate of the cost. Namely, we need to find the coefficient of n^3 in the polynomial giving the total number of operations. As usual, the terms in n^2 and n are ignored.

In terms of computational cost by iteration, the Bregman splitting Algorithm 3.1 is, in general, the most expensive while Algorithm 3.2 is the cheapest. However, Algorithm 3.1 converges, in general, faster, attaining the same accuracy of the other two algorithms in much fewer iterations (see Section 5).

Algorithm 3.1: The cost is mainly determined by the cost of solving a Sylvester-type equation of the form (21) and the computation of two Singular Value Decompositions (SVD). Solving (21) involves about $76n^3$ operations ($66n^3$ for the QZ algorithm and $10n^3$ for the remaining calculations; see [34]), and each SVD involves about $22n^3$ operations by the method of Golub and Reinsch [16]. Note that while each iteration includes two SVD, the quite expensive QZ decomposition is just required one time because A_1 and A_2 are fixed during all the iterations.

Algorithms 3.2 and 3.3: The objective functions considered in the steepest descent-type algorithms involve the computation of the trace of matrix products. The efficient computation of $\text{trace}(AB)$ does not require matrix-matrix products. Instead, it can be carried out through the formula:

$$\text{trace}(AB) = \sum_{i,j} (A \circ B^T)_{(i,j)}, \quad (24)$$

where the operator \circ denotes the Hadamard product, i.e., the entry-wise product. If A and B are matrices of order n , the direct computation of the matrix product AB needs $O(n^3)$ operations, while the trace at (24) just requires $O(n^2)$. However, as far as we know, the trace

of a product of three or four matrices requires the computation of one matrix-matrix product which costs $2n^3$ operations.

Hence, the evaluation of the objective function \tilde{F} (see its expression in (6)) involved in the steepest descent type algorithms requires $4n^3$ operations (two matrix-matrix products), while its computation inside the cycles requires only $2n^3$; the product JM needs to be computed just one time as the algorithm runs.

Each iteration of the Algorithm 3.2 requires the computation of a Cayley transform, which corresponds to solving a multiple right-hand side linear system of the form $AX = B$, which costs about $8n^3/3$. Concerning Algorithm 3.3, one exponential of a skew-symmetric matrix is required in each iteration. In [9], a scaling and squaring algorithm designed specifically for exponentials of a skew-symmetric matrix is proposed, with an overall cost of $2(16/3 + s)n^3$, where s stands for the number of squarings. Alternatively, one can use the general algorithm available through the function `expm` of MATLAB, which implements the scaling and squaring algorithm of [3]. Its cost is $O(n^3)$, and we refer the reader to [3] for the detailed expression of the computational cost. We note that, in the particular case $n = 3$, the exponential of a skew-symmetric matrix can be computed by the well-known Rodrigues's formula, at the cost of just one matrix-matrix product.

4.3. Residual Estimates

Let us consider the residual function

$$R(X) := XJ - JX^T - M \tag{25}$$

and assume that \tilde{X} is an approximation to the exact solution X of the Moser-Veselov equation obtained by a certain numerical algorithm. Hence, $\tilde{X} = X + \Delta$, for some matrix Δ of order n .

In the numerical computations of solutions of matrix equations, it is, in general, difficult to estimate the absolute error $\|\Delta\| = \|X - \tilde{X}\|$ or the relative error $\|\Delta\|/\|X\|$, where $\|\cdot\|$ stands for a given subordinate matrix norm. Thus, the authors work instead with relative residuals to check the quality of the approximation \tilde{X} and the numerical stability.

An obvious definition for the relative residual of the Moser-Veselov equation would be

$$\|R(\tilde{X})\|/\|X\|. \tag{26}$$

However, as pointed out in [17, Sec. 5] (see also [18, Problem 7.15]) for the matrix equation $X^p = A$ (i.e., for computing the matrix p th root of A), this definition may not be appropriate in some situations. This has motivated us to propose a definition for the relative residual in the style of what is suggested in [17, 18].

With respect to the residual function given in (25), we have

$$\begin{aligned} R(\tilde{X}) &= (X + \Delta)J - J(X + \Delta)^T - M \\ &= XJ + \Delta J - JX^T - J\Delta^T - M \\ &= \Delta J - J\Delta^T. \end{aligned}$$

By vectorization and attending to some properties of the Kronecker product, we have

$$\begin{aligned} \text{vec}\left(R(\tilde{X})\right) &= \text{vec}(\Delta J) - \text{vec}(J\Delta^T) \\ &= (J \otimes I) \text{vec}(\Delta) - (I \otimes J) \text{vec}(\Delta^T) \\ &= (J \otimes I - (I \otimes J)\Pi) \text{vec}(\Delta), \end{aligned}$$

where Π is the permutation matrix of order $n^2 \times n^2$ (check (18)). Denoting $C := J \otimes I - (I \otimes J)\Pi \in \mathbb{R}^{n^2 \times n^2}$, we have

$$\text{vec}\left(R(\tilde{X})\right) = C \text{vec}(\Delta).$$

With respect to the spectral norm $\|\cdot\|_2$, we have

$$\left\| \text{vec}\left(R(\tilde{X})\right) \right\|_2 \leq \|C\|_2 \|\text{vec}(\Delta)\|_2,$$

and, by attending that, for any matrix A , $\|\text{vec}(A)\|_2 = \|A\|_F$, it follows

$$\left\| R(\tilde{X}) \right\|_F \leq \|C\|_2 \|\Delta\|_F. \quad (27)$$

Let us suppose that $\|\Delta\|_F \leq \epsilon \|X\|_F$, for a certain small value ϵ . Note that $\|Q\|_F = \sqrt{n}$, for any orthogonal matrix Q of order n because $\|Q\|_F^2 = \text{trace}(Q^T Q) = \text{trace}(I) = n$. Then $\|\Delta\|_F \leq \epsilon \sqrt{n}$ and

$$\|R(\tilde{X})\|_F \leq \epsilon \sqrt{n} \|C\|_2,$$

or, equivalently,

$$\frac{\|R(\tilde{X})\|_F}{\sqrt{n} \|C\|_2} \leq \epsilon.$$

This suggests the following definition for the relative residual:

$$\rho(\tilde{X}) := \frac{\|R(\tilde{X})\|_F}{\sqrt{n} \|C\|_2}. \quad (28)$$

To understand why the relative residual (28) is more meaningful than (26), we shall notice that solving $XJ - JX^T = M$ (ignoring the orthogonal constraint on X) is equivalent to solve the linear system $C \text{vec}(X) = \text{vec}(M)$, with C being in general singular, and that the norm of C must be considered in error analysis, as the expression (28) does.

4.4. Termination Criteria

Several strategies are available to decide when terminating an iterative procedure. Attending to the nature of our optimization problem in (4), the sequence generated by $F(X_k)$, where

$$F(X) = \|XJ - JX^T - M\|_F^2,$$

must converge to zero. Hence, we can fix a tolerance ϵ and then iterate while $F(X_k) \geq \epsilon$. Note that, in Algorithms 3.2 and 3.3, $F(X_k)$ is available during the algorithm and does not involve extra cost.

The typical behavior of methods with linear convergence is to slow down as X_k approaches stationary points. Often, it may be difficult to detect this phenomenon's occurrence, so another condition to stop the cycle must be added. In our implementations of the algorithm, we consider the classical relative difference

$$\frac{\|X_k - X_{k-1}\|_F}{\|X_k\|_F} = \frac{\|X_k - X_{k-1}\|_F}{\sqrt{n}}$$

as well. Because the algorithms are implemented in finite precision environments, a maximal number of iterations must be considered to stop iterating. In summary, we fix tolerances ϵ_1 , ϵ_2 and a maximal number of iterations k_0 and stop iterating when

$$\frac{\|X_k - X_{k-1}\|_F}{\sqrt{n}} < \epsilon_2 \quad \text{or} \quad k > k_0. \quad (29)$$

5. Numerical Experiments

To evaluate the performance of the proposed algorithms, we have carried out a set of experiments in MATLAB R2021a (with unit roundoff $u \approx 1.1 \times 10^{-16}$) in a machine with

Core i5 (1.60GHz). To terminate the iteration procedure in the algorithms, we have used the following criteria:

$$\|X_k - X_{k-1}\|_F / \sqrt{n} \leq \text{tol} \quad \text{or} \quad k > 1000,$$

where `tol` is a prescribed tolerance. The following terminology is used:

- `#iter`: number of iterations;
- `rel-res`: relative residual defined in (28);
- $\tilde{F}(X)$: value of the objective function defined in (6);
- $\|\text{grad } \tilde{F}(X)\|_F$: norm of the Riemannian gradient defined in (13).

We have considered $r = 1$ in Algorithm 3.1 and $\tau = 10^{-3}$ in Algorithm 3.2. Most of the Moser-Veselov equations considered in the experiments do not satisfy the condition $M^2/4 + J^2 > 0$ required by direct methods, but the associated Hamiltonian matrix \mathcal{H} (see (3)) has no pure imaginary eigenvalue. Experiment 4 is devoted to the case where \mathcal{H} has pure imaginary eigenvalues.

5.1. Experiment 1

This experiment involves a set of 100 Moser-Veselov equations with randomized matrices J and M of order 16, plus a set of 100 Moser-Veselov equations with randomized matrices J and M of order 17, and so on, up to a set of 100 Moser-Veselov equations with randomized matrices J and M of order 35 and aims at comparing Algorithm 3.1 with Algorithm 3.2 in terms of `#iter`, `rel-res`, $\tilde{F}(X)$ and $\|\text{grad } \tilde{F}(X)\|_F$, for a tolerance `tol` = 10^{-10} and the initial guess set to $X_0 = I$. The results are displayed through a boxplot with whiskers in Figures 1 and 2. Experiment 1 involves 2000 Moser-Veselov equations and, as expected, there may have outliers, which are represented by red crosses in the graphs. All of them are finite, i.e., no NaN's or Inf's arose in the calculations. In most of the cases, these outliers reflect difficulties in the intermediate calculations involved in the algorithms, like the computation of inverses of ill-conditioned.

A careful inspection of those figures lead us to conclude that Algorithm 3.1 gives the best results in terms of relative residuals, number of iterations, values of the objective function and

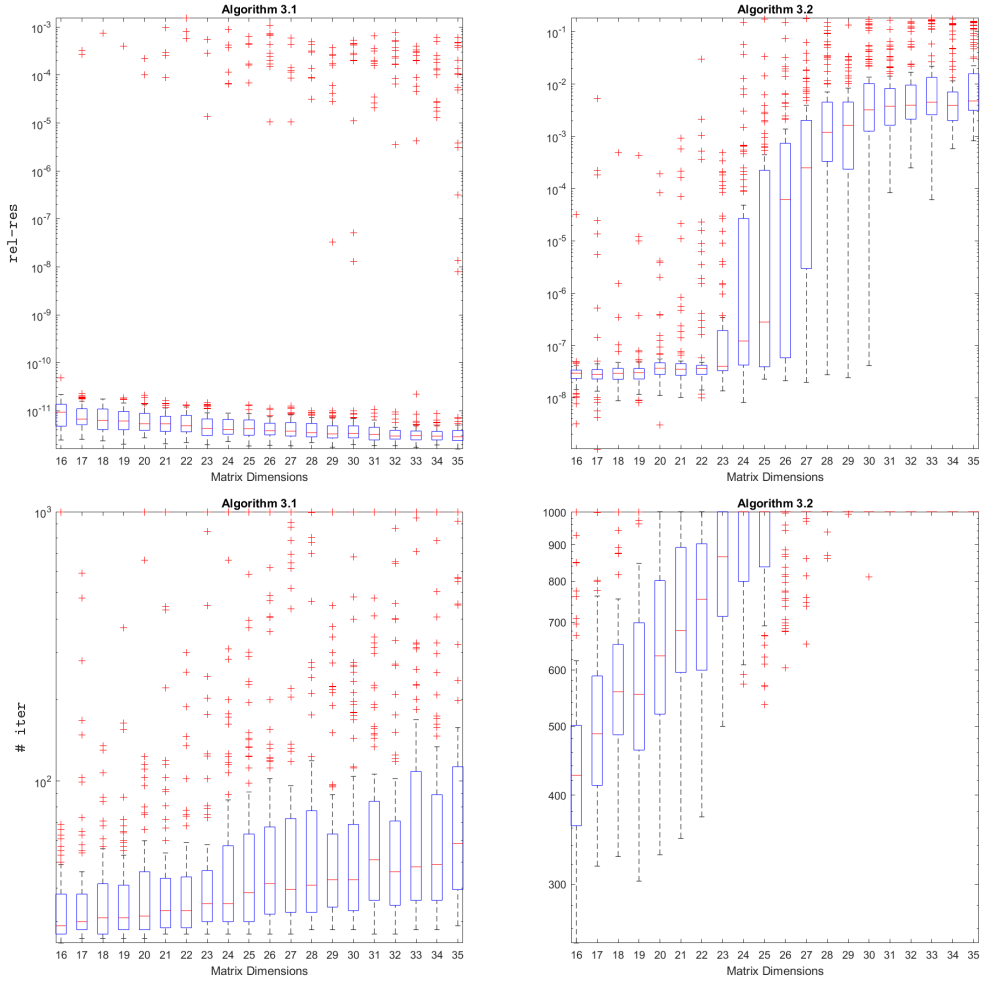


Figure 1: Comparison between Algorithm 3.1 and Algorithm 3.2 in terms of $\#iter$ and $rel-res$, for a tolerance $tol = 10^{-10}$ and $X_0 = I$.

norms of the Riemannian gradient, despite it involves higher computational cost by iteration than Algorithm 3.2. However, since it requires much fewer iterations, the overall computational cost is, in general, smaller.

5.2. Experiment 2

This experiment involves 1000 Moser-Veselov equations: 100 equations with randomized matrices J and M of order 6, plus 100 equations with randomized matrices J and M of order

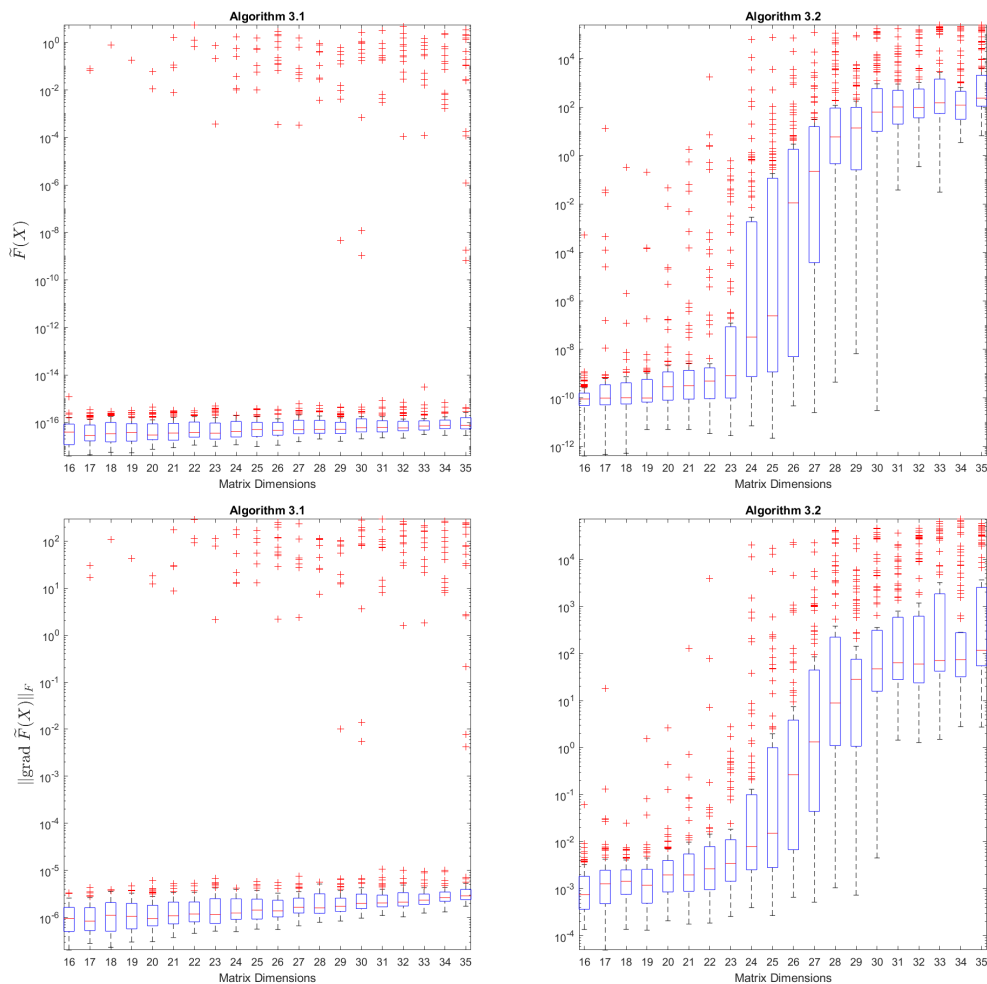


Figure 2: Comparison between Algorithm 3.1 and Algorithm 3.2 in terms of $\tilde{F}(X)$ and $\|\text{grad } \tilde{F}(X)\|_F$, for a tolerance $\text{tol} = 10^{-10}$ and $X_0 = I$.

7, and so on, up to 100 equations with randomized matrices J and M of order 15. It illustrates the performance of Algorithms 3.1 and 3.2 according to the choice of the initial guess X_0 , in terms of the number of iterations `#iter` and of relative residuals `rel-res`. The results of Algorithm 3.1 are displayed in Figure 3 and of Algorithm 3.2 in Figure 4, by means of boxplot graphs with whiskers. We observe that both algorithms perform much better if we take $X_0 = I$ instead of a randomized special orthogonal matrix.

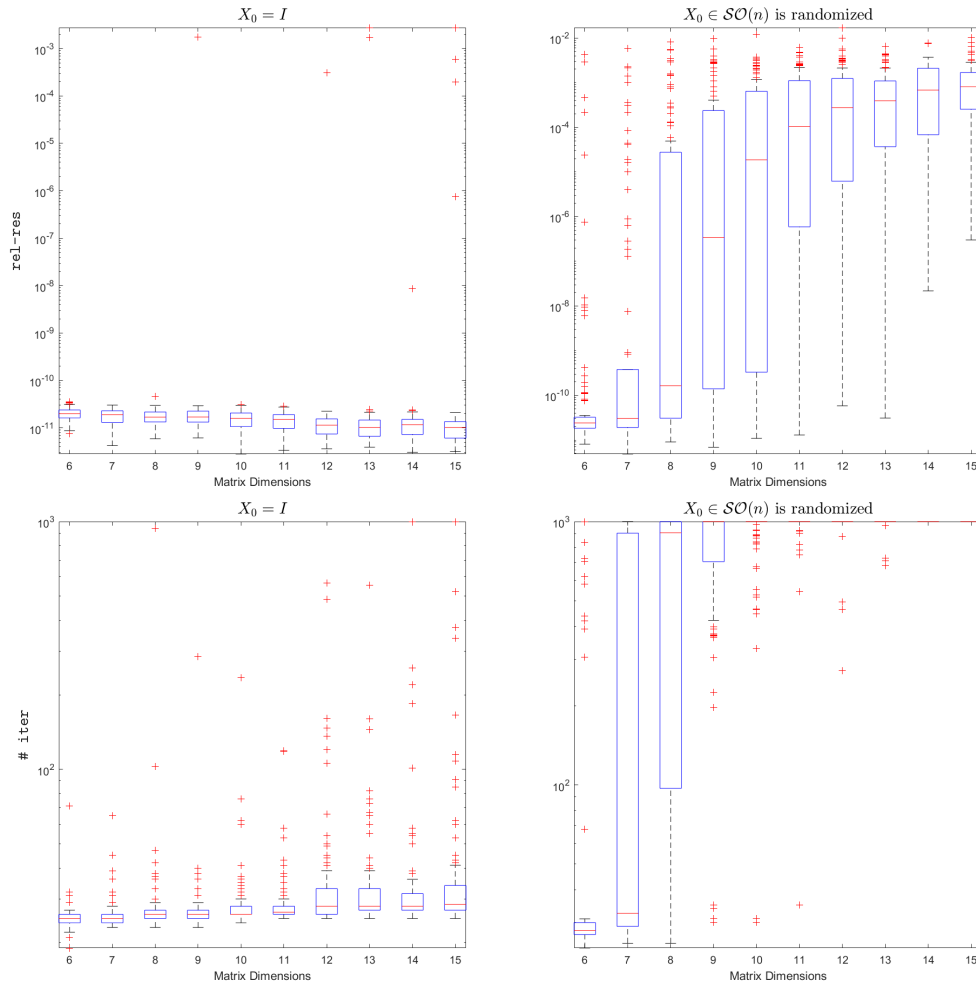


Figure 3: Relative residual and number of iterations of Algorithm 3.1 according to the choice of the starting approximation X_0 for a tolerance $\text{tol} = 10^{-10}$.

5.3. Experiment 3

Now, we consider the same set of 1000 Moser-Veselov equations as in Experiment 2 (Sec. 5.2) to illustrate Algorithm 3.1, SOC Algorithm + `fminunc` (i.e., the algorithm described at the beginning of Sec. 3.1, with Step 2(a) solved using `fminunc` of MATLAB, where the gradient (16) is specified in the options mode), Algorithm 3.2 and Algorithm 3.3, in terms of relative residuals and the number of iterations, for $X_0 = I$ and a tolerance $\text{tol} = 10^{-5}$. The results are displayed in Figure 5. Note that now we are using a larger tolerance than in Experiments

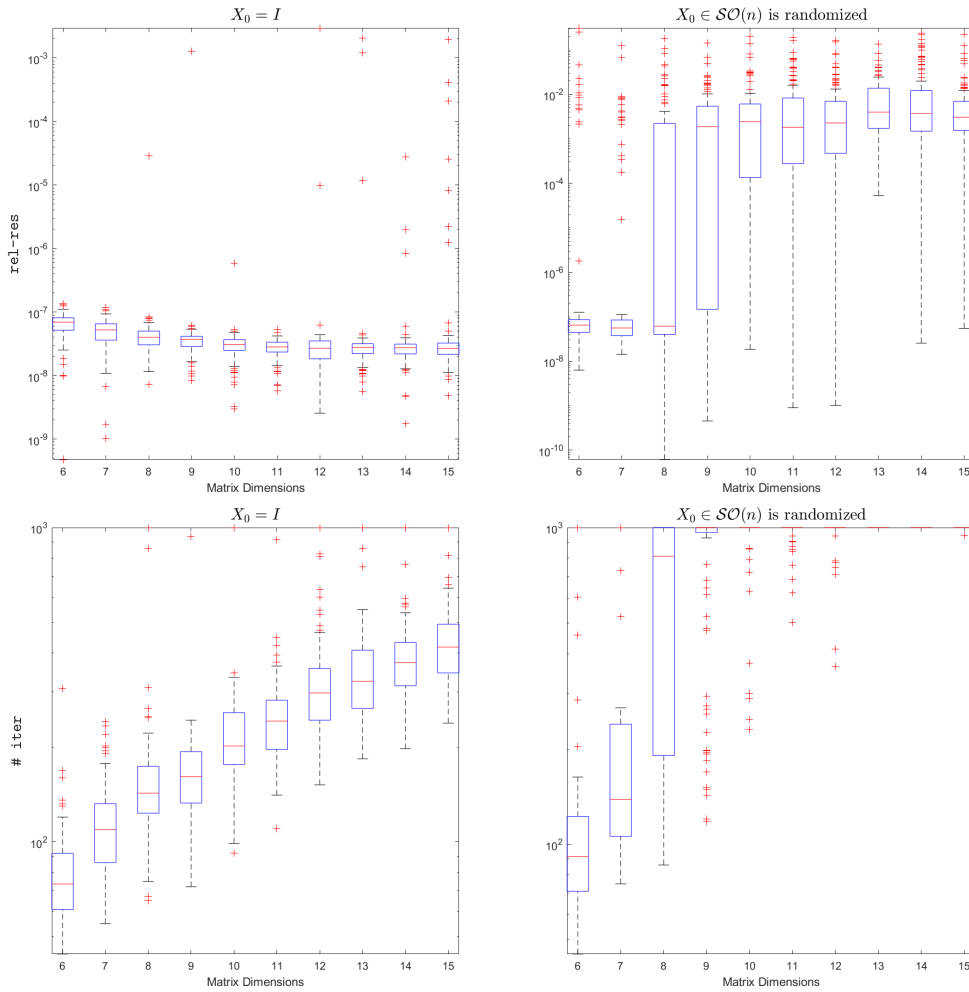


Figure 4: Relative residual and number of iterations of Algorithm 3.2 according to the choice of the starting approximation X_0 for a tolerance $\text{tol} = 10^{-10}$.

1 and 2. This is because for smaller tolerances both SOC Algorithm+`fminunc` and Algorithm 3.3 diverge frequently or may require thousands of iterations.

In Figure 6, we compare the computational time of Algorithm 3.1 with SOC Algorithm + `fminunc`, where we can see that the latter algorithm requires about 100 times the computational time of the former. If we decrease the tolerance or increase the size of the matrices, the results of SOC Algorithm + `fminunc` get even worse and may have little practical interest. For these cases, we recommend instead Algorithms 3.1 or 3.2.

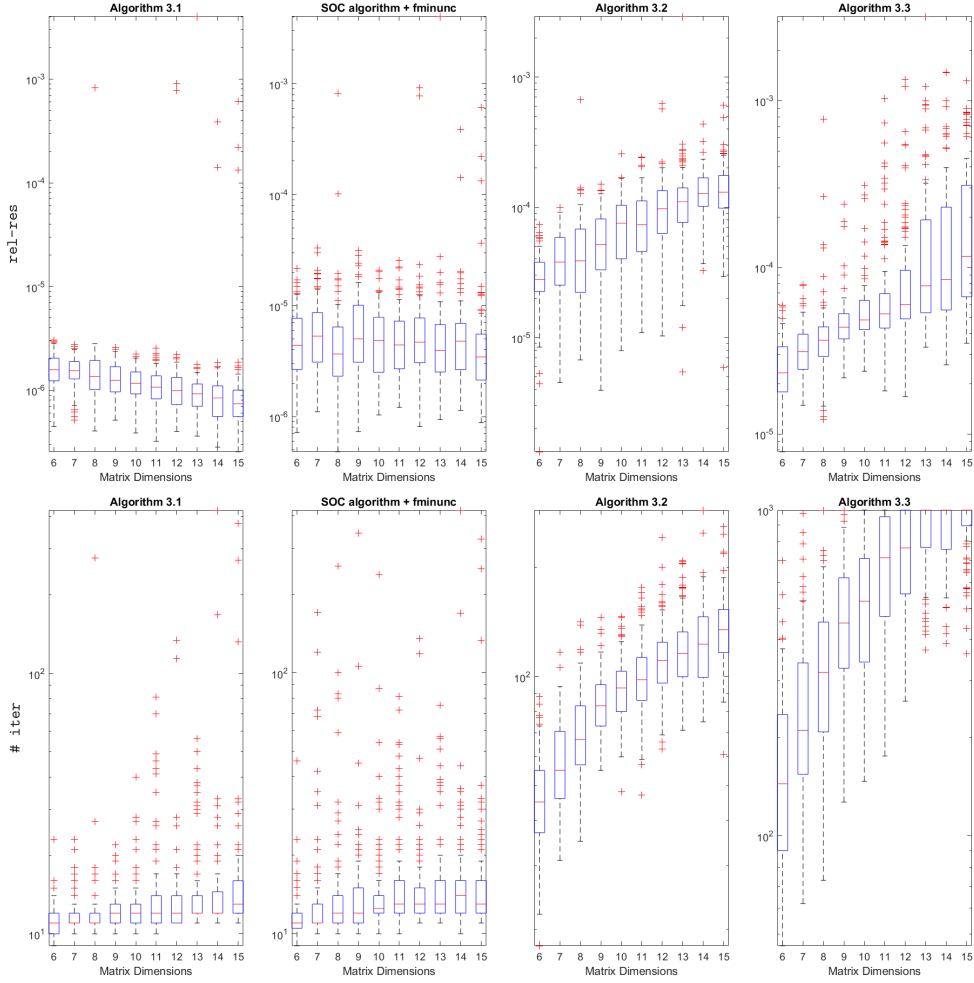


Figure 5: Relative residuals and number of iterations of Algorithm 3.1, SOC Algorithm + *fminunc*, Algorithm 3.2 and Algorithm 3.3, for $X_0 = I$ and $\text{tol} = 10^{-5}$.

5.4. Experiment 4

In this experiment, the goal is to illustrate the behaviour of Algorithms 3.1, 3.2 and 3.3 when the Hamiltonian matrix \mathcal{H} (check (3)) has some pure imaginary eigenvalues associated to Jordan blocks of even sizes. We have taken ten Moser-Veselov equations involving randomized matrices of order 4 that were carefully chosen to guarantee that \mathcal{H} fits the conditions just mentioned. The results are displayed in Figure 7. While in Experiments 1, 2 and 3, Algorithm 3.1 has performed very well in terms of the number of iterations, now it is the one that gives the




Figure 6: Comparison between the computational time of Algorithms 3.2 and $SOC+fminunc$, for $X_0 = I$ and $\tau_{ol} = 10^{-5}$.

poorest results (1000 iterations in all the ten tests)! For a compromise regarding the relative residual and the number of iterations, Algorithm 3.2 gives the best results.

5.5. Other experiments and considerations

Experiments with Algorithms 3.1 and 3.2 for Moser-Veselov equations involving randomized dense matrices of larger sizes ($n = 80, 90, 100, 150, 200$) were also performed (not shown here), for a tolerance $\tau_{ol} = 10^{-10}$. In terms of the magnitude of the relative residual and of computational time, we observed a gradual deterioration as the size of the matrices increased. So, our overall recommendation is that those algorithms are suitable for dense matrices of small/medium size, say $n \leq 100$.

We are not aware of existing algorithms for solving Moser-Veselov equations involving dense (or sparse) matrices of large size. We leave these investigations for further research.

Methods for solving the optimization problem in (4) hardly give results with relative residuals of order the unit roundoff $u \approx 1.1 \times 10^{-16}$. That is the price to pay for considering the objective function as $F(X) = \|XJ - JX^T - M\|_F^2$ instead of $\|XJ - JX^T - M\|_F$, whose expression is more difficult to handle. In iterative methods implemented in environments involving floating-point arithmetic, as $F(X)$ becomes less than u , its value may stop decreasing.

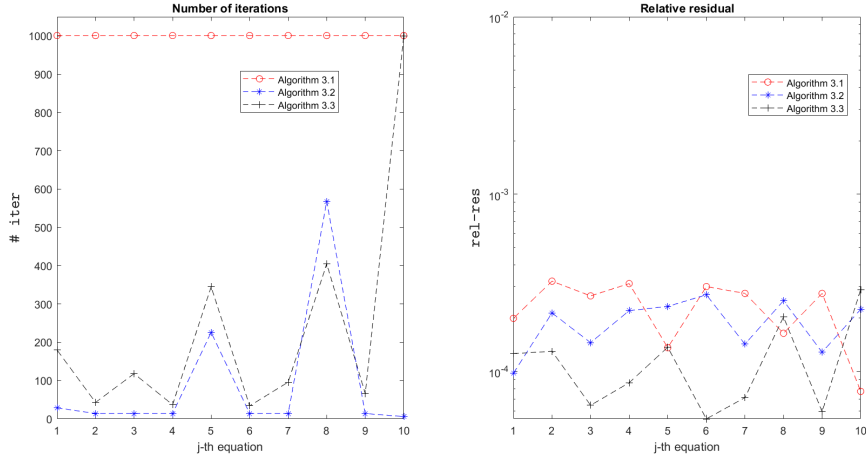


Figure 7: Results for Algorithms 3.1, 3.2 and 3.3 when \mathcal{H} in (3) has some pure imaginary eigenvalues associated to Jordan blocks of even sizes; $X_0 = I$ and $\text{tol} = 10^{-5}$.

So, it is more reasonable to expect approximations with relative residuals of order \sqrt{u} .

6. Conclusions

This paper proposes three algorithms for solving the discrete Euler-Arnold equation for the generalized rigid body motion estimation: a Bregman splitting (Algorithm 3.1), and two steepest descent-based methods (Algorithms 3.2 and 3.3). An essential advantage of these methods is that they do not require the strong condition $M^2/4 + J^2 > 0$, in contrast with other methods existing in the literature. Important numerical issues related to the algorithms, like convergence, computational cost, residual estimates, and termination criteria, have been investigated in detail. The numerical experiments that have been carried out to evaluate the performance of the three methods suggest that the Bregman splitting Algorithm 3.1 is promising, at least for equations where the associated Hamiltonian matrix has no pure imaginary eigenvalue.

Acknowledgements: The work of João R. Cardoso was partially supported by the Centre for Mathematics of the University of Coimbra - UIDB/00324/2020, funded by the Portuguese Government through FCT/MCTES. Pedro Miraldo was partially supported by the LARSyS - FCT Plurianual funding 2020-2023.

References

References

- [1] Abrudan, T., Eriksson, J., Koivunen, V., 2008. Steepest descent algorithms for optimization under unitary matrix constraint. *Trans. Sig. Proc.* 56, 1134–1147.
- [2] Absil, P.A., Mahony, R., Sepulchre, R., 2007. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press.
- [3] Al-Mohy, A., Higham, N., 2009. A new scaling and squaring algorithm for the matrix exponential. *SIAM J. Matrix Anal. Appl.* 31, 970–989.
- [4] Barzilai, J., Borwein, J., 1988. Two-point step size gradient methods. *IMA J. Numer. Anal.* 8, 141–148.
- [5] Bloch, A., 2015. *Nonholonomic Mechanics and Control*. 2 ed., Springer.
- [6] Bregman, L., 1967. The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex optimization. *USSR Comput. Math. Math. Phys.* 7, 200–217.
- [7] Campos, J., Cardoso, J.R., Miraldo, P., 2019. Poseamm: A unified framework for solving pose problems using an alternating minimization method, in: *IEEE Int’l Conf. Robotics and Automation (ICRA)*, pp. 3493–3499.
- [8] Cardoso, J., Leite, F., 2003. The Moser-Veselov equation. *Linear Alg. Appl.* 360, 237–248.
- [9] Cardoso, J., Leite, F., 2010. Exponentials of skew-symmetric matrices and logarithms of orthogonal matrices. *J. Comput. Appl. Math.* 233, 2867–2875.
- [10] Coleman, T., Li, Y., 1994. On the convergence of reflective newton methods for large-scale nonlinear minimization subject to bounds. *Mathematical Programming* 67, 189–224.
- [11] Coleman, T., Li, Y., 1996. An interior, trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization* 6, 418–445.

- [12] Dai, Y.H., Fletcher, R., 2005. Projected barzilai-borwein methods for large-scale box-constrained quadratic programming. *Numer. Math.* 100, 21–47.
- [13] Dang, Q., Gui, H., Liu, K., Zhu, B., 2020. Relaxed-constraint pinpoint lunar landing using geometric mechanics and model predictive control. *Journal of Guidance, Control, and Dynamics* URL: <https://doi.org/10.2514/1.G005039>.
- [14] Edelman, A., Arias, T., Smith, S., 1999. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.* 20, 303–353.
- [15] Gao, B., Liu, X., Chen, X., Yuan, Y.X., 2018. A new first-order algorithmic framework for optimization problems with orthogonality constraints. *SIAM J. Optim.* 28, 302–332.
- [16] Golub, G., Reinsch, C., 1970. Singular value decomposition and least squares solutions. *Numer. Math.* 14, 403–420.
- [17] Guo, C.H., Higham, N., 2006. A schur-newton method for the matrix p-th root and its inverse. *SIAM J. Matrix Anal. Appl.* 28, 788–804.
- [18] Higham, N., 2008. *Functions of Matrices: Theory and Computation*. SIAM, Philadelphia, PA, USA.
- [19] Jiang, B., Dai, Y.H., 2015. A framework of constraint preserving update schemes for optimization on stiefel manifold. *Math. Program.* 153, 535–575.
- [20] Kalabic, U., Gupta, R., Cairano, S.D., Bloch, A.M., Kolmanovsky, I.V., 2014. Constrained spacecraft attitude control on $SO(3)$ using reference governors and nonlinear model predictive control. *American Control Conference* , 5586–5593.
- [21] Kalabic, U., Gupta, R., Cairano, S.D., Bloch, A.M., Kolmanovsky, I.V., 2017. Mpc on manifolds with an application to the control of spacecraft attitude on $SO(3)$. *Automatica* 76, 293–300.
- [22] Lai, R., Osher, S., 2014. A splitting method for orthogonality constrained problems. *J. Sci. Comput.* 58, 431–449.

- [23] Lee, T., Leok, M., McClamroch, N., 2007. Lie group variational integrators for the full body problem in orbital mechanics. *Celestial. Mech. Dyn. Astr.* 98, 121–144.
- [24] Luenberger, D., Ye, Y., 2015. *Linear and Nonlinear Programming*. Springer Publishing Company, Incorporated.
- [25] Lutkepohl, H., 1996. *Handbook of Matrices*. 1 ed., Jonh Wiley and Sons.
- [26] Manton, J., 2002. Optimization algorithms exploiting unitary constraints. *Trans. Sig. Proc.* 50, 635–650.
- [27] McLachlan, R., Zanna, A., 2005. The discrete moser—veselov algorithm for the free rigid body, revisited. *Found. Comput. Math.* 5, 87–123.
- [28] Miraldo, P., Cardoso, J.R., 2020. On the generalized essential matrix correction: An efficient solution to the problem and its applications. *J. Math. Imaging Vis.* 62, 1107–1120.
- [29] Moser, J., Veselov, A.P., 1991. Discrete versions of some classical integrable systems and factorization of matrix polynomials. *Commun. Math. Phys.* 139, 217–243.
- [30] Nocedal, J., Wright, S., 2006. *Numerical Optimization*. second ed., Springer, New York, NY, USA.
- [31] Nordkvist, N., Sanyal, A., 2010. A lie group variational integrator for rigid body motion in SE(3) with applications to underwater vehicle dynamics. *IEEE Conference on Decision and Control (CDC)* , 5414–5419.
- [32] Osher, S., Burger, M., Goldfarb, D., Xu, J., Yin, W., 2005. An iterative regularization method for total variation-based image restoration. *Multiscale Model. Simul.* 4, 460–489.
- [33] Polak, E., 1997. *Optimization: Algorithms and Consistent Approximations*. Springer-Verlag.
- [34] Terán, F.D., Dopico, F., 2011. Consistency and efficient solution of the sylvester equation for \star -congruence: $AX + X^*B = C$. *Electron.J. Linear Algebra* 22, 849–863.

- [35] Wen, Z., Yin, W., 2013. A feasible method for optimization with orthogonality constraints. *Math. Program.* 142, 397–434.
- [36] Zhu, X., 2017. A riemannian conjugate gradient method for optimization on the stiefel manifold. *Comp. Opt. and Appl.* 67, 73–110.