SE(3) **Synchronization by eigenvectors of dual quaternion matrices**

Ido Hadi [ID]*

*School of Mathematical Sciences, Faculty of Exact Sciences, Tel Aviv University, Tel Aviv 6997801,
Israel*
*Corresponding author. Email: idohadi@mail.tau.ac.il

Tamir Bendory [ID]

School of Electrical Engineering, Faculty of Engineering, Tel Aviv University, Tel Aviv 6997801, Israel

AND

Nir Sharon

*School of Mathematical Sciences, Faculty of Exact Sciences, Tel Aviv University, Tel Aviv 6997801,
Israel*

In synchronization problems, the goal is to estimate elements of a group from noisy measurements of their ratios. A popular estimation method for synchronization is the spectral method. It extracts the group elements from eigenvectors of a block matrix formed from the measurements. The eigenvectors must be projected, or 'rounded', onto the group. The rounding procedures are constructed ad hoc and increasingly so when applied to synchronization problems over non-compact groups. In this paper, we develop a spectral approach to synchronization over the non-compact group SE(3), the group of rigid motions of $\mathbb{R}^3$. We based our method on embedding SE(3) into the algebra of dual quaternions, which has deep algebraic connections with the group SE(3). These connections suggest a natural rounding procedure considerably more straightforward than the current state of the art for spectral SE(3) synchronization, which uses a matrix embedding of SE(3). We show by numerical experiments that our approach yields comparable results with the current state of the art in SE(3) synchronization via the spectral method. Thus, our approach reaps the benefits of the dual quaternion embedding of SE(3) while yielding estimators of similar quality.

Keywords: applied non-commutative algebra; dual quaternions; group synchronization; spectral algorithms.

1. Introduction

Synchronization problems arise as part of data processing pipelines in several contexts, including single-particle reconstruction in cryogenic electron microscopy (3; 21), structure from motion problems (2; 27) and simultaneous localization and mapping problems (1; 23). A synchronization problem is an estimation problems over a group G, in which group elements $g_1, \ldots, g_n \in G$ are estimated from measurements of their ratios $g_i g_j^{-1}$. These measurements are inherently ambiguous, since the set of ratios of $g_1 g, \ldots, g_n g \in G$, for any $g \in G$, is the same as the set of ratios of the original group elements. Thus, in synchronization problems, the goal is to estimate the group elements up to a right-multiplication by an arbitrary element of G.

The spectral method is a widely used method for solving synchronization problems. It seeks to approximate a solution by finding eigenvectors of a block matrix formed by embedding the group into a matrix algebra. The embedding enables the richer algebraic structure of matrix algebras, and in particular their spectral decomposition, to be leveraged to solve the synchronization problem. Importantly, the eigenvectors of the observation matrix are not themselves a solution of the synchronization problem. They have to be projected back onto the group, a projection that is typically non-linear. This procedure is often referred to as 'rounding' in the synchronization literature. In typical rounding procedures, a matrix formed by these eigenvectors is chopped into blocks and these blocks are subsequently mapped onto elements of the group.

Consider a synchronization problem over a relatively simple compact group in the absence of measurement noise. As we explain in greater detail in Section 3.2, even in this ideal setting, typically there is a gap between the synchronization problem and the eigenproblem used to solve it. Oversimplifying to an extent, we can say that every solution to synchronization problem is a solution of the eigenproblem used by the spectral method, but not every solution of the eigenproblem is a solution of the synchronization problem. This issue is exacerbated for synchronization over non-compact groups, the focal point of this paper. The solutions of the eigenproblem form a compact set, the set of unit eigenvectors which span the relevant eigenspaces of the measurement matrix. Yet, unlike the case of synchronization over compact groups, the possible solutions of a synchronization problems over non-compact groups form a non-compact set. Therefore, the gap between the two problems can be said to be much larger.

In practice, this gap between the eigenproblem and the synchronization problem must be bridged by the rounding procedure. It ensures that the output of the synchronization method is comprised of elements of the group. However, existing rounding procedures are made to do so in an ad hoc manner. Specifically, they do not stem from a well-characterized relationship between the synchronization problem and the eigenproblem or the relationship between their respective algebraic contexts, the group and the algebra in which it is embedded. They are merely constructed to ensure that the result of a synchronization problem is an element of the group.

In this paper, we address the issues we raised above in the non-compact case of SE(3) synchronization, the group of rigid motions, rotations and translations, of $\mathbb{R}^3$. In (1), an analog of the spectral method was developed for SE(3) synchronization, which utilized an embedding of SE(3) into a matrix algebra. We replace this matrix algebra embedding with an embedding into the algebra of dual quaternions. Unlike the matrix algebra embedding used by (1), the algebra of dual quaternions is much closer to the structure of SE(3), in a sense we explain in Section 2.4. Two recent advances in dual quaternion linear algebra were a spectral theorem for matrices of dual quaternions (17) and a power iteration capable of approximating its spectra (7). These two results allow us to develop an elegant spectral synchronization method. The algebraic context eliminates the gap between the synchronization problem and the eigenproblem, at least in the absence of measurement noise. In addition, the rounding step itself stems directly from the algebraic relationship between elements of the algebra itself and the set on which SE(3) is represented. In addition, it is computationally simpler than the rounding procedure of (1).

The structure of this paper is as follows. We begin in Section 2 with a survey of the necessary background material. Among other things, we define the algebra of dual quaternions, matrix algebras with dual quaternion entries and their properties. In Section 3, we delve deeper into synchronization problems and the spectral method. We define the problem rigorously, introduce the spectral method and conclude by laying out our approach, relying on dual quaternion embedding of SE(3). Finally, in Section 4, we demonstrate empirically that our approach yields comparable results with the current state-of-the-art spectral method for SE(3) synchronization developed by (1). Thus, we offer a better theoretical

foundation tying SE(3) synchronization problems with eigenproblems while still maintaining the quality of estimation of the current state-of-the-art spectral method.

2. Background: dual quaternions and matrices of dual quaternions

We provide a succinct background information on dual quaternions and dual quaternion matrices. In Section 2.1 and Section 2.2, we construct the dual quaternions and describe some of their properties. In Section 2.3 we survey the well-known representation of SO(3) on the algebra of quaternions and in Section 2.4 we survey the lesser known representation of SE(3) on the algebra of dual quaternions. In Section 2.5 we define the free modules of dual quaternions and algebras of matrices with dual quaternion elements. This section culminates in the spectral theorem for matrices of dual quaternions, which was recently proved by (17). Finally, in Section 2.6 we survey the power iteration for matrices of dual quaternions, which was recently developed by (7).

2.1 Dual numbers

The algebra of dual numbers is the set $\mathbb{D} = \{a + b\boldsymbol{\varepsilon} \mid a, b \in \mathbb{R}\}$ with addition and multiplication defined by

$$\left(a_1 + b_1\boldsymbol{\varepsilon}\right) + \left(a_2 + b_2\boldsymbol{\varepsilon}\right) = \left(a_1 + a_2\right) + \left(b_1 + b_2\right)\boldsymbol{\varepsilon} \tag{2.1}$$

$$\left(a_1 + b_1\boldsymbol{\varepsilon}\right)\left(a_2 + b_2\boldsymbol{\varepsilon}\right) = a_1 a_2 + \left(a_1 b_2 + b_1 a_2\right)\boldsymbol{\varepsilon}. \tag{2.2}$$

If $a + b\boldsymbol{\varepsilon} \in \mathbb{D}$, it is convenient to refer to a as the real coordinate and b as the dual coordinate. It is easy to see from their definition that the dual numbers are defined analogously to the complex numbers. Although the complex numbers are the algebra over the reals generated by 1 and i such that $i^2 = -1$, the algebra of dual numbers is generated by 1 and $\boldsymbol{\varepsilon}$ such that $\boldsymbol{\varepsilon}^2 = 0$. This similarity accounts for the similar properties of the algebras of dual numbers and complex numbers. Namely, the underlying vector space of both is $\mathbb{R}^2$, addition and multiplication operations of both are associative and commutative and both algebras have a unit, 1.

Despite the aforementioned similarities, these two algebras are fundamentally very different. The algebra of complex numbers is a field, while the algebra of dual numbers is not. It has zero divisors, i.e. non-zero elements whose product is zero. Indeed, recall that an element x of an algebra is said to be nilpotent if there is a positive integer n such that $x^n = 0$. The only nilpotent complex number is 0. The following proposition characterizes the nilpotent dual numbers, and it is obvious it has many non-zero nilpotents.

PROPOSITION 1. An $x = a + b\boldsymbol{\varepsilon} \in \mathbb{D}$ is nilpotent if and only if $a = 0$. Also, $x^2 = 0$ for every nilpotent.

Proof. If $a \neq 0$, it is easy to show that (2.2) implies that $x^n = a^n + (\ldots)\boldsymbol{\varepsilon} \neq 0$. If $a = 0$, then $x^2 = b^2\boldsymbol{\varepsilon}^2 = 0$. □

As the next proposition shows, nilpotent dual numbers are the only dual numbers which are not invertible.

PROPOSITION 2. If $x = a + b\boldsymbol{\varepsilon} \in \mathbb{D}$, x is invertible if and only it is not nilpotent. In that case, $x^{-1} = a^{-1} - ba^{-2}\boldsymbol{\varepsilon}$.

Proof. Let $y = c + d\boldsymbol{\varepsilon}$. From (2.2) it follows that $xy = 1$ if

$$ac = 1, \qquad ad + bc = 0.$$

The first equation has a solution if and only if $a \neq 0$. By proposition 1, this holds if and only if x is not nilpotent. The formula for x^{-1} is easily obtained by solving this system of equations when $a \neq 0$. □

A final point of analogy between the dual and complex numbers is that both admit a square root. Following (18), given $x = a + b\varepsilon$ and $y = c + d\varepsilon$, we write $x > y$ if $a > c$, or $a = c$ and $b > d$. We say a dual number $x = a + b\varepsilon \in \mathbb{D}$ is **non-negative** if $x \geq 0$ and **positive** if $x > 0$. In (18, Sec. 2), the following definition of a square root was given:

$$\sqrt{x} = \begin{cases} \sqrt{a} + \frac{b}{2\sqrt{a}}\varepsilon & x \text{ is positive and invertible,} \\ 0 & x = 0. \end{cases} \tag{2.3}$$

In all other cases, the square root is undefined. Using this definition, it is possible to show that $\sqrt{xy} = \sqrt{x}\sqrt{y}$. This multiplicative identity implies that for any $x = a + b\varepsilon \in \mathbb{D}$ with $a \neq 0$, we have $\sqrt{x^2} = |a| + \text{sgn}(a)b\varepsilon$, where $\text{sgn}(a) = 1$ for $a > 0$, zero when $a = 0$ and -1 when $a < 0$. Such considerations motivated (18) to define the following absolute value on dual numbers:

$$|x| = \begin{cases} |a| + \text{sgn}(a)b\varepsilon & a \neq 0 \\ |b|\,\varepsilon & a = 0. \end{cases}$$

It has the following properties (18, Theorem 2), which are generalized forms of the usual absolute value defined on the real line:

THEOREM 3. For any $x, y \in \mathbb{D}$:

1. $|x| = 0$ if and only if $x = 0$.

2. $|x| = x$ if $x \geq 0$ and $|x| > x$ otherwise.

3. $|x| = \sqrt{x^2}$ if x is invertible.

4. $|xy| = |x|\,|y|$ for any $x, y \in \mathbb{D}$.

5. $|x + y| \leq |x| + |y|$.

2.2 Dual quaternions

The algebra of dual quaternions is the easiest to define using the algebra of quaternions. The algebra of quaternions is the set $\mathbb{H} = \{a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k} \mid a, b, c, d \in \mathbb{R}\}$. We typically refer to b, c and d as the $\mathbf{i}$, $\mathbf{j}$ or $\mathbf{k}$ coordinates, respectively. In the literature, a is sometimes referred to as the real part of the quaternion, but in order to eliminate confusion with the real and dual parts of a dual number, we simply refer to it as the first coordinate of the quaternion. Addition in $\mathbb{H}$ is defined coordinate-wise and multiplication is defined by the following relations among its generators:

$$\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1. \tag{2.4}$$

As a vector space, $\mathbb{H} \cong \mathbb{R}^4$. Also, $\mathbb{R} \subset \mathbb{H}$ by constraining the $\mathbf{i}$, $\mathbf{j}$ and $\mathbf{k}$ coordinates to zero. The centre of algebra, the set of elements for which multiplication commutes with all other elements of the algebra, is $\mathbb{R}$. There is an involution on the quaternions, defined analogously to the complex conjugation, $q^* = a - b\mathbf{i} - c\mathbf{j} - d\mathbf{k}$ for $q = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$. This involution is referred to as **quaternion conjugation** and

we say that q^* is the **conjugate** of q. Similarly to the complex norm, there is a norm on the quaternions defined as $|q| := \sqrt{q^*q} = \sqrt{a^2 + b^2 + c^2 + d^2}$, which can be shown to imply $\mathbb{H} \cong \mathbb{R}^4$ as a normed vector space. We note here that $|q|^2 = q^*q = qq^*$, and so the inverse of every non-zero $q \in \mathbb{H}$ is $\frac{q^*}{q^*q}$. Overall, the algebra of quaternions is an involutive normed division algebra.

The algebra of dual quaternions is $\mathbb{DH} = \{a + b\boldsymbol{\varepsilon} \mid a, b \in \mathbb{H}\}$. Essentially, it means a dual quaternion is a dual number with quaternion coordinates instead of real coordinates. From this, we immediately see that $\mathbb{D} \subset \mathbb{DH}$ and also that $\mathbb{H} \subset \mathbb{DH}$, the former by constraining $a, b \in \mathbb{R}$ and the latter by setting $b = 0$. Addition and multiplication are still defined by (2.1) and (2.2) with the quaternion binary operations replacing the real ones. The trivial involution on $\mathbb{D}$, the identity and quaternion conjugation induce an involution on $\mathbb{DH}$, $x^* = a^* + b^*\boldsymbol{\varepsilon}$ for $x = a + b\boldsymbol{\varepsilon} \in \mathbb{DH}$. Because $\mathbb{D}$ is the centre of $\mathbb{DH}$, we can use the usual notation for division whenever $d = a + b\boldsymbol{\varepsilon} \in \mathbb{D}$ is invertible and $x \in \mathbb{DH}$, and so we write

$$\frac{x}{d} := xd^{-1} = d^{-1}x. \tag{2.5}$$

As a vector space, $\mathbb{DH} \cong \mathbb{R}^8$. We use the same names for the coordinates we used for the dual numbers. Given $a + b\boldsymbol{\varepsilon} \in \mathbb{DH}$, we call a the real coordinate and b the dual coordinate.

The algebra of dual numbers is a unital algebra, since both $\mathbb{D}$ and $\mathbb{H}$ are. It is non-commutative, since $\mathbb{H}$ is. It is not a division algebra, since $\mathbb{D}$ is not. This combination of features sets it apart from both the algebra of quaternions and the algebra of dual numbers. Despite this, many of their features generalize well to the dual quaternions. As we did for the dual numbers, we begin by noting that its nilpotents are characterized in exactly the same way as the nilpotents of the dual numbers. Indeed, one needs to only replace $\mathbb{D}$ with $\mathbb{DH}$ in proposition 1 to obtain the proper characterization.

Note here that zero is the only nilpotent in the algebra of quaternions. Therefore, one can say the first part of proposition 2 holds for quaternions, namely, that a quaternion is invertible if and only if it is not nilpotent. We show now that the dual quaternions share this property. Before that, we prove a lemma pointing out several important features of the involutive structure of $\mathbb{DH}$.

LEMMA 4. If $x = a + b\boldsymbol{\varepsilon} \in \mathbb{DH}$, then $x^*x = |a|^2 + (ab^* + ba^*)\boldsymbol{\varepsilon}$ is a dual number and $x^*x = xx^*$.

Proof. The first equality is easily proved by substituting x^* and x into (2.2). To prove the result is a dual number, we note first that the norm of a quaternion is a real number. In addition, $q^* + q \in \mathbb{R}$ for every quaternion q and $(ab^*)^* = ba^*$, which implies $ab^* + ba^*$ is also real. The second identity follows from a direct calculation. $\square$

PROPOSITION 5. If $x = a + b\boldsymbol{\varepsilon} \in \mathbb{DH}$, x is invertible if and only if it is not nilpotent. In that case, $x^{-1} = \frac{x^*}{x^*x}$.

Proof. If $a = 0$, we have $xy = bc\boldsymbol{\varepsilon} \neq 1$ for every $y = c + d\boldsymbol{\varepsilon} \in \mathbb{DH}$. Therefore, x is invertible only if $a \neq 0$, that is, only if x is not nilpotent. If $a \neq 0$, it follows from Lemma 4 that the real coordinate of x^*x is positive. Thus, by proposition 2 x^*x is invertible. Finally, taking $y = \frac{x^*}{x^*x}$ one obtains that $yx = 1$. The equality $xy = 1$ follows easily from the second equality in Lemma 4. $\square$

In (18), absolute value was also defined for dual quaternions. If $x = a + b\boldsymbol{\varepsilon} \in \mathbb{DH}$, then

$$|x| = \begin{cases} |a| + \frac{ab^* + ba^*}{2|a|}\boldsymbol{\varepsilon} & a \neq 0 \\ |b|\,\boldsymbol{\varepsilon} & a = 0. \end{cases} \tag{2.6}$$

It can be readily seen that when x is a dual number, (2.6) and (2.3) agree. In (18, Thm. 4), the following properties of the absolute value of dual quaternions were proved:

THEOREM 6. Let $x, y \in \mathbb{DH}$.

1. $|x|$ is a dual number.

2. If x is invertible, then $|x| = \sqrt{qq^*}$.

3. $|x| = |x^*|$.

4. $|x| \geq 0$ for all x and $|x| = 0$ if and only if $x = 0$.

5. $|xy| = |x|\,|y|$.

6. $|x + y| \leq |x| + |y|$.

2.3   *Representing* SO(3) *on the quaternions*

The algebra of quaternions, which we surveyed in Section 2.2, has a close relationship with the special orthogonal group SO(3), the group of norm and orientation-preserving linear transformations of the Euclidean space $\mathbb{R}^3$. The subset $\mathbb{H}_1 := \{q \in \mathbb{H} \mid |q| = 1\}$ is a subgroup of the multiplicative group of $\mathbb{H}$. This follows easily from the fact $|q|^2 = q^*q = qq^*$. In particular, this also implies that the inverse of every $q \in \mathbb{H}_1$ is its conjugate q^*, making the group structure of $\mathbb{H}_1$ closed under involution and dependent on the involutive structure of the algebra of quaternions. A well-known result states that $\mathbb{H}_1$ is a double cover of SO(3) (8, Sec. 4.9). We succinctly construct the covering map here. We begin by identifying $\mathbb{R}^3 \cong \mathbb{T}$ as vector spaces, where $\mathbb{T} = \{a\mathbf{i} + b\mathbf{j} + c\mathbf{k} \mid a, b, c \in \mathbb{R}\} \subset \mathbb{H}$. Then, for every $q \in \mathbb{H}_1$, define a map $\mathbb{R}^3 \to \mathbb{R}^3$ by $\varphi(q)(x) = qxq^*$, where $x \in \mathbb{R}^3$ is treated as a quaternion in $\mathbb{T}$ via the identification above. We then have the following:

PROPOSITION 7. The covering map is a 2-to-1 surjective group homomorphism $\varphi : \mathbb{H}_1 \to$ SO(3), such that $\varphi^{-1}(g) = \{q, -q\}$ for every $g \in$ SO(3).

This double cover allows one to represent elements of SO(3) on $\mathbb{H}_1$. More specifically, one can represent SO(3) on $\mathbb{H}_{1+} = \{q \in \mathbb{H}_1 \mid q + q^* \geq 0\}$ by choosing an appropriate $q \in \varphi^{-1}(g)$.

2.4   *Representing* SE(3) *on the dual quaternions*

The **special Euclidean group** SE(3) is the group of rigid motions of $\mathbb{R}^3$, i.e. rotations and translations of $\mathbb{R}^3$, but not reflections. It is a semidirect product SE(3) = SO(3) $\ltimes \mathbb{R}^3$. Its elements are most directly represented as pairs of the form $(R, t) \in$ SO(3) $\times \mathbb{R}^3$ with the group operation defined as $(R_1, t_1) \circ (R_2, t_2) = (R_1 R_2, t_1 + R_1 t_2)$. Here, $\circ$ can be thought of as the composition of two affine transformations.

Let $\mathbb{DH}_1 = \{x \in \mathbb{DH} \mid x^*x = 1\}$ be the set of **unit dual quaternions**. Let $\mathbb{DT} = \left\{1 + \frac{1}{2}t\boldsymbol{\varepsilon} \mid t \in \mathbb{T}\right\}$, where we defined $\mathbb{T}$ in Section 2.3. Note that $\mathbb{H}_1 \subset \mathbb{DH}_1$ and $\mathbb{DT} \subset \mathbb{DH}_1$. We state an analogue of Proposition 7. It is a synopsis of (20, Chp. 9) and its proof can be found there.

PROPOSITION 8. We have

1. $\mathbb{DH}_1 = \mathbb{DT} \rtimes \mathbb{H}_1$.

2. Let $\psi : \mathbb{DH}_1 \to$ SE(3) be defined by $\psi(q, t)(x) = (\varphi(q), t)$. Here, φ is the double cover defined in proposition 7, $q \in \mathbb{H}_1$ and $1 + \frac{1}{2}t\boldsymbol{\varepsilon} \in \mathbb{DT}$ so t is identified with an element of $\mathbb{R}^3$. Then ψ is a 2-to-1 surjective group homomorphism such that $\psi^{-1}(g) = \{-x, x\}$ for all $g \in$ SE(3).

Proposition 8 offers a way to represent elements of SE(3) using dual quaternions in much the same way we represent elements of SO(3) on the quaternions. We represent $(R, \mathbf{t}) \in$ SE(3) by $x = qt \in \mathbb{DH}_1$, where q is a unit quaternion with non-negative first coordinate representing the rotation R and $\mathbf{t}$ is represented by $t = 1 + \frac{1}{2}t'\boldsymbol{\varepsilon}$, $t' = t_1\mathbf{i} + t_2\mathbf{j} + t_3\mathbf{k}$. Its action on $\mathbb{R}^3$ can be defined similarly to what we saw in the quaternion case. Given $\mathbf{s} \in \mathbb{R}^3$, let $s = 1 + s'\boldsymbol{\varepsilon}$, $s' = t_1\mathbf{i} + t_2\mathbf{j} + t_3\mathbf{k}$. Note the absence of a $\frac{1}{2}$ factor in this embedding of $\mathbb{R}^3$ into the dual quaternions, compared with how the translational part is represented on $\mathbb{DT}$ in proposition 8. The action of $(R, \mathbf{t}) \in$ SE(3) on $\mathbf{s}$ is then $s \mapsto xsx^*$. Indeed, we have

$$
\begin{aligned}
xsx^* &= \left(q + \frac{1}{2}t'q\boldsymbol{\varepsilon}\right)\left(1 + s'\boldsymbol{\varepsilon}\right)\left(q^* + \frac{1}{2}q^*t'^*\boldsymbol{\varepsilon}\right) \\
&= qq^* + \left(\frac{1}{2}t'qq^* + qs'q^* + \frac{1}{2}qq^*t'^*\right)\boldsymbol{\varepsilon} \\
&= 1 + \left(qs'q^* + t'\right)\boldsymbol{\varepsilon}.
\end{aligned}
$$

In expanding the products in the second transition we used the fact $\boldsymbol{\varepsilon}^2 = 0$, which implied that no products of two dual coordinates could appear in the result. The dual coordinate of the result is exactly the action of $(R, \mathbf{t})$ on $\mathbf{s}$, because $qs'q^*$ is the action of R on $\mathbf{s}$ by proposition 7 and because of the manner in which we represented $\mathbb{R}^3$ on quaternions when we constructed t' and s'.

In (7), it was proved that it is possible to project an almost arbitrary dual quaternion onto the unit dual. Their result is stated here as follows:

THEOREM 9. Let $x = a + b\boldsymbol{\varepsilon} \in \mathbb{DH}$ be non-zero.

1. If $a \neq 0$, then $\frac{q}{|q|}$ is a unit dual quaternion solving

$$
\min_{v \in \mathbb{DH}_1} |v - x|^2. \tag{2.7}
$$

2. If $a = 0$, then any $q' = a' + b'\boldsymbol{\varepsilon}$ such that $a' = \frac{b}{|b|}$ and $b^*b' + b'^*b = 0$ solves (2.7).

This theorem exposes an aspect of the close tie between SE(3) and the algebra of dual quaternions. We focus on its first part. It can be interpreted as the idea that every invertible dual quaternion is essentially a unit dual quaternion multiplied by a dual number. Thus, in a sense, the multiplicative structure of the algebra of dual quaternions is almost entirely generated by the group of unit dual quaternions. This property has a geometric interpretation via the absolute value function. Namely, dividing a dual quaternion by its absolute value yields a solution to the optimization problem (2.7), a sort of projection function. This function itself is defined via the dual quaternion absolute value (2.6), which itself intimately related to the algebraic structures on the algebra of dual quaternions as indicated in Theorem 6.

Finally, we refer the interested reader to Appendix A, where we compare various representations of SE(3), focusing on features which are perhaps most salient in applied work.

2.5 *Matrices of dual quaternions and their spectral decomposition*

Modules are an extensively studied generalization of the notion of a vector space to vectors and matrices with entries from an arbitrary algebra. See (10) for an introduction to module theory. Here, we provide a synopsis of the necessary information. A (right-)module $\mathbb{M}$ over an algebra $\mathbb{A}$ is a commutative group

$(\mathbb{M}, +)$ along with a map $\mathbb{M} \times \mathbb{A} \to \mathbb{M}$, $(m, x) \mapsto m \cdot x$, which satisfies the following properties. First, it is distributive in the sense that $m \cdot (x + y) = m \cdot x + m \cdot y$ and $(m_1 + m_2) \cdot x = m_1 \cdot x + m_2 \cdot x$. Second, the operation is associative, $(m \cdot x) \cdot y = m \cdot (x \cdot y)$. We often refer to this map as the action of $\mathbb{A}$ on $\mathbb{M}$ and omit the multiplication sign. Oversimplifying to an extent, a module can be thought of as a vector space in which the underlying field was replaced by an arbitrary algebra.

We focus on two types of free modules of the algebra of dual quaternions, $\mathbb{DH}^n$ and $\mathbb{DH}^{n \times n}$. Given a positive integer n, the **free $\mathbb{DH}$-module** is $\mathbb{DH}^n = \left\{ \mathbf{m} = (m_1, \ldots, m_n)^\top \mid m_j \in \mathbb{DH}, j = 1, \ldots, n \right\}$. It is acted upon by $\mathbb{DH}$ via entry-wise right-multiplication, i.e. $(\mathbf{m}x)_j = m_j x$. Addition is defined entry-wise. Moreover, $\mathbb{DH}^{n \times n}$ can be thought of as the module of all $n \times n$ matrices of dual quaternions. As an $\mathbb{DH}$-module, one can identify it with $\mathbb{DH}^{n^2}$, and so addition is also performed entry-wise and so is right-multiplication by a dual quaternion. However, it actually is an algebra, not just a module, with the multiplication of matrices defined in the familiar way. Similarly, one can define the product $\mathbf{Ax}$ of a matrix $\mathbf{A} \in \mathbb{DH}^{n \times n}$ and $\mathbf{x} \in \mathbb{DH}^n$ in the usual way. It is instructive to think of $\mathbf{A}$ as a $\mathbb{DH}$-linear operator on $\mathbb{DH}^n$ in the sense that $\mathbf{A}(\mathbf{x}a + \mathbf{y}b) = (\mathbf{Ax})a + (\mathbf{Ay})b$ for any $\mathbf{x}, \mathbf{y} \in \mathbb{DH}^n$ and $a, b \in \mathbb{DH}$.

Several types of real and complex matrices have dual quaternions counterparts. The conjugate transpose of a matrix $\mathbf{A} \in \mathbb{DH}^{n \times n}$ is $(\mathbf{A}^*)_{ij} = \mathbf{A}_{ji}^*$. This can easily be shown to be an involution on $\mathbb{DH}^{n \times n}$. We say that a matrix $\mathbf{A} \in \mathbb{DH}^{n \times n}$ is Hermitian if $\mathbf{A} = \mathbf{A}^*$. It is unitary if $\mathbf{A}^* \mathbf{A} = \mathbf{A} \mathbf{A}^* = \mathbf{I}_n$, where $\mathbf{I}_n$ is the identity matrix, which is defined in usual way.

Deepening the analogy with real and complex matrices, a notion of eigenvector and eigenvalue exists for matrices of dual quaternions. We say $\lambda \in \mathbb{DH}$ is a (right-)eigenvalue of the matrix $\mathbf{A} \in \mathbb{DH}^{n \times n}$ if there is $\mathbf{x} \in \mathbb{DH}^n$ such that $\mathbf{Ax} = \mathbf{x}\lambda$. As the name suggests, the non-commutativity of $\mathbb{DH}$ means one can also define a left-eigenvalue of a matrix of dual quaternions, though we will not be presenting any results on these. Whenever we speak of eigenvalues and eigenvectors of dual quaternion matrices, we will be referring to right-eigenvalues and right-eigenvectors.

We are now ready to state the spectral theorem for Hermitian matrices of dual quaternions (17, Thm. 4.1):

THEOREM 10. Let $\mathbf{A} \in \mathbb{DH}^{n \times n}$. If $\mathbf{A}$ is Hermitian, then there are a unitary $\mathbf{U} \in \mathbb{DH}^{n \times n}$ and a diagonal $\boldsymbol{\Sigma} \in \mathbb{DH}^{n \times n}$ with dual number elements such that $\mathbf{A} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{U}^*$. The diagonal elements of $\boldsymbol{\Sigma}$ have the form $\sigma_j = \alpha_j + \beta_j \boldsymbol{\varepsilon} \in \mathbb{D}$ with $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_n$. Here, σ_j is the eigenvalue of $\mathbf{A}$ corresponding to the eigenvector $\mathbf{u}_j$, the jth column of $\mathbf{U}$.

This theorem essentially states that Hermitian matrices of dual quaternions have a full set of eigenvectors. Furthermore, in the dual quaternion case, as in the complex case, the eigenvalues turn out to be self-adjoint, that is, they satisfy $\sigma^* = \sigma$, with respect to the appropriate involution. This highlights the striking similarities between this theorem and the well-known spectral theorem of complex Hermitian matrices.

2.6 *Power iteration for hermitian matrices of dual quaternion*

In a recent paper, (7) developed a dual quaternion analog of the power iteration, capable of approximating the top eigenvalue of a Hermitian dual quaternion matrix and its corresponding eigenvalue. We clarify below the sense in which an eigenvalue is the top eigenvalue. Let $\mathbf{A} \in \mathbb{DH}^{n \times n}$ be a Hermitian matrix of dual quaternions. Given some initializer $\mathbf{v}^0 \in \mathbb{DH}^n$, the power iteration defined in (7) is

$$\mathbf{y}^{(k)} = \mathbf{A}\mathbf{v}^{(k-1)}, \quad \lambda^{(k-1)} = \left(\mathbf{v}^{(k-1)}\right)^* \mathbf{y}^{(k)}, \quad \mathbf{v}^{(k)} = \frac{\mathbf{y}^{(k)}}{\left\| \mathbf{y}^{(k)} \right\|_2}. \tag{2.8}$$

Here, the dual quaternion norm used is defined as follows:

$$\|\mathbf{x}\|_2 = \begin{cases} \sqrt{\sum_{j=1}^n \left|x_j'\right|^2} \boldsymbol{\varepsilon} & \mathbf{x} = \mathbf{x}'\boldsymbol{\varepsilon}, \mathbf{x}' \in \mathbb{H}^n, \\ \sqrt{\sum_{j=1}^n \left|x_j\right|^2} & \text{otherwise.} \end{cases} \tag{2.9}$$

For simplicity, we refer to $\|\mathbf{x}\|_2$ as a norm, even though it is not a norm in the usual linear-algebraic sense. Several additional definitions are required. Following (7), we write $c_k = \tilde{O}_D\left(s^k\right)$ for some $s < 1$ if $c_k = a_k + b_k \boldsymbol{\varepsilon} \in \mathbb{D}$ and for some polynomial $h(k)$ we have $a_k = O\left(s^k h(k)\right)$ and $b_k = O\left(s^k h(k)\right)$. The following result was proved in (7):

THEOREM 11. (Theorem 4.1 of (7))Let $\mathbf{u}_1, \ldots, \mathbf{u}_n$ be the eigenvectors of $\mathbf{A}$ ordered so that their corresponding eigenvalues $\sigma_j = \alpha_j + \beta_j \boldsymbol{\varepsilon} \in \mathbb{D}$ satisfy $\left|\alpha_1\right| > \left|\alpha_2\right| \geq \cdots \geq \left|\alpha_n\right|$. If $\mathbf{v}^{(0)} = \sum_{j=1}^n \mathbf{u}_j \gamma_j$ with $\gamma_j \in \mathbb{DH}$ such that γ_1 has a non-zero real coordinate, then the sequence defined by (2.8) satisfies

$$\mathbf{v}^{(k)} = \mathbf{u}_1 c \left(1 + \tilde{O}_D\left(\left|\frac{\alpha_2}{\alpha_1}\right|^k\right)\right), \quad \lambda^{(k)} = \sigma_1 \left(1 + \tilde{O}_D\left(\left|\frac{\alpha_2}{\alpha_1}\right|^{2k}\right)\right),$$

where $c = \mathrm{sgn}(\alpha_1)\frac{\gamma_1}{|\gamma_1|}$.

This theorem essentially shows that the power iteration (2.8) converges to $\mathbf{u}_1$ and its corresponding eigenvalue. At every iteration, $\mathbf{v}^{(k)}$ and $\lambda^{(k)}$ are essentially the desired values multiplied by some dual number, which approaches 1 at an exponential rate. In this and in the general form of the iteration (2.8), the dual quaternion power iteration is strikingly similar to the well-established power iteration for real and complex matrices (25, Lecture 27).

2.7 Isometries of eigenspaces of matrices of dual quaternions

At this point, we have established the background material and notation required to point out two important facts. First, the eigenspaces of a Hermitian matrix of dual quaternion are invariant to the action of unit dual quaternions. Indeed, if $x \in \mathbb{DH}$ and $\mathbf{x} \in \mathbb{DH}^n$ is an eigenvector of $\mathbf{A} \in \mathbb{DH}^{n \times n}$ corresponding to eigenvalue $\lambda \in \mathbb{D}$, then

$$\mathbf{A}(\mathbf{x}x) = \mathbf{x}\lambda x = (\mathbf{x}x)\lambda.$$

Therefore, $\mathbf{x}x$ is also an eigenvector corresponding to eigenvalue λ. The last transition follows, because x and λ commute, since the latter is in the centre of $\mathbb{DH}$.

The second fact relates to the isometries of $\mathbb{DH}^n$ with respect to the dual quaternion norm defined in (2.9). Given $\mathbf{x} \in \mathbb{DH}^n$ and an invertible $x \in \mathbb{DH}$, we have

$$\|\mathbf{x}x\|_2 = \|\mathbf{x}\|_2 |x|.$$

This can be proved by considering both cases of (2.9) separately and using Theorem 6. Therefore, $\|\mathbf{x}\| = \|\mathbf{x}x\|_2$ if and only $|x| = 1$, which is equivalent to $x^*x = 1$, i.e. x is a unit dual quaternion. Therefore, the

unit dual quaternions are the isometries of a free module of dual quaternions with respect to the norm structure induced by (2.9).

Combining these two facts together, we conclude that for a fixed unit dual number q, the $\mathbb{D}\mathbb{H} \to \mathbb{D}\mathbb{H}$ map $\mathbf{x} \mapsto \mathbf{x}q$ has two important properties. First, it preserves the norm (2.9). Second, it maps eigenvectors of Hermitian matrices of dual quaternions to eigenvectors corresponding to the same eigenvalue. These two properties play a key role in the spectral approach to synchronization we develop in the next section.

3. The spectral approach to synchronization problems via dual quaternions

In this section, we apply the dual quaternion representation of SE(3) to the problem of SE(3) synchronization. In Section 3.1, we describe group synchronization problems and noise types which were considered in the literature. In Section 3.2, we describe the spectral method, a prevalent solution to group synchronization problems. As we describe in Section 3.3, this approach can be generalized to the dual quaternions algebra and therefore can be used to address SE(3) synchronization problems. We dedicate Section 4 to numerical experiments demonstrating that our proposed method indeed works on simulated data.

3.1 *Group synchronization: problem statement*

Let G be a group. Let $g_1, \ldots, g_n$ be elements of G and let $g_{ij} = g_i g_j^{-1}$ be their ratios, $1 \le i,j \le n$. A synchronization problem is an estimation problem in which one attempts to estimate $\{g_i \mid 1 \le i \le n\}$, the absolute group elements, from their ratios $\{g_{ij} \mid 1 \le i < j \le n\}$. Often, the clean ratios are perturbed by some sort of noise. Because $g_1 g, \ldots, g_n g$ generate the same set of clean measurements for every arbitrary $g \in G$, we can only hope to estimate $g_1, \ldots, g_n$ up to a global right-multiplication by an arbitrary element of G.

In order to describe the typical noise models, it is convenient to represent the group faithfully as some subgroup of the general linear group, the group of invertible matrices. Thus, we assume in this section that one has an injective homomorphism $\rho : G \to \mathrm{GL}_k(\mathbb{F})$, with $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$. The absolute group elements and their ratios are mapped onto $\mathrm{GL}_k(\mathbb{F})$ via ρ in a consistent manner. Namely, $\rho(g_{ij}) = \rho(g_i)\rho(g_j)^{-1}$. The elements of G are represented by elements of $\mathbb{F}^{k \times k}$, the $k \times k$ matrices with entries in $\mathbb{F}$. We refer to $\mathbb{F}^{k \times k}$ as the representation space and to the image of ρ as the representing subspace. Throughout this subsection, we suppress the homomorphism ρ and assume the absolute group elements and their ratios are already faithfully represented in some matrix algebra.

Given a group represented in this manner, three kinds of noise sources are usually considered in the literature. First, **perturbative noise** can be applied to each clean ratio. This has two forms. **Multiplicative noise** perturbs the clean ratio while still remaining in the group, so that we measure $\widehat{g}_{ij} = g_{ij}\varepsilon_{i,j}$, with $\varepsilon_{i,j} \in G$. **Additive noise** perturbs the clean ratio within the representation space, so that we measure $\widehat{g}_{ij} = g_{ij} + \varepsilon_{i,j}$, with $\varepsilon_{i,j} \in \mathbb{F}^{k \times k}$. We here assume that perturbative noise is either multiplicative for all ratios, or additive for all ratios.

The last two kinds of noise are best viewed as applied to a set of measurements. This set of measurements may be assumed to be either a complete set of clean measurements $\{g_{ij} \mid 1 \le i < j \le n\}$, a measurement set afflicted with perturbative noise, $\{\widetilde{g}_{ij} \mid 1 \le i < j \le n\}$, or a subset of one of these. Every set of measurements can be envisaged as an undirected graph of order n, in which each vertex is labeled with an absolute group element and the edges correspond to the measurements in the set. See Fig. 1 for an example.
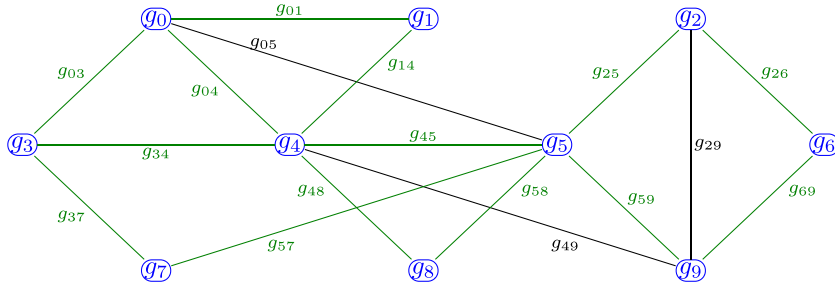
FIG. 1. An example of a synchronization measurement graph with 10 absolute group elements, labeled from 0 to 9. Vertices appear in blue circles. Edges are labeled with ratio measurements. Edges of the subgraph of non-corrupted measurements are dyed green. Edge labels of this subgraph can be thought of as either clean or afflicted with perturbative noise. In a synchronization problem, the goal is to estimate the absolute group elements in a manner that is as consistent as possible with the measurements of their ratios.

The first kind of noise acting on the measurement graph, which we term **selection noise**, merely removes a subset of the edges of the measurement graph. The second, which is often called **corruptive** or **adversarial noise**, replaces the labels of a subset of edges with a randomly generated labels which are independent from the clean group ratios. In both kinds of noise, the affected subset of edges may be chosen randomly or deterministically. The measurement graph has a subgraph formed by edges which were not corrupted. We assume that this subgraph is connected, regardless of the method used to choose the subset of edges affected by selection or corruptive noise. Figure 1 shows an example of a measurement graph with corruptive noise.

All combinations of the three kinds of noise were considered in the literature, both in numerical investigations of estimation methods and in their theoretical analysis (1; 5; 6; 12; 13; 19; 22; 26). The literature cited provides a mere illustrative selection.

3.2 *Group synchronization via the spectral method*

Let $g_1, \ldots, g_n$ be elements of a group G. We emphasize at the outset of this subsection that all elements of G are faithfully represented within some matrix algebra $\mathbb{F}^{k \times k}$. Whenever perturbative noise is considered, it is always assumed to be multiplicative.

The spectral method relates synchronization problems to eigenproblems on a matrix, which is essentially a generalized adjacency matrix of the synchronization graph. Let $\mathbf{Y} \in \mathbb{F}^{nk \times nk}$ be a block matrix such that the (i,j)th block $\mathbf{Y}_{ij}$, $i < j$, is the measurement of the ratio $g_i g_j^{-1}$. If the measurement is clean, perturbed (multiplicatively) or corrupted, we assign $\mathbf{Y}_{ji} = \mathbf{Y}_{ij}^{-1}$. If the measurement is missing, we set $\mathbf{Y}_{ij} = \mathbf{Y}_{ji} = \mathbf{0}_{k \times k}$. The diagonal blocks are $\mathbf{Y}_{ii} = \mathbf{I}_k$ for $i = 1, \ldots, n$. Essentially, the block on the ith row of blocks and jth column of blocks is a measurement of the ratio of $g_i g_j^{-1}$. Since these measurements label the edges of the synchronization graph, this is a generalization of the notion of the adjacency matrix of a graph.

In the absence of noise, it is evident that $\mathbf{Y} = \mathbf{g}\mathbf{g}^{-1}$. Here, $\mathbf{g} \in \mathbb{F}^{nk \times k}$ is the column block matrix with g_i as its ith block and $\mathbf{g}^{-1} \in \mathbb{F}^{k \times nk}$ is the row block matrix with g_i^{-1} as its ith block. It is immediate that $\mathbf{Y}\mathbf{g} = n\mathbf{g}$ and that n is the only non-zero eigenvalue of $\mathbf{Y}$. Thus, the columns of $\mathbf{g}$ form a basis of the eigenspace of $\mathbf{Y}$ associated with n. Furthermore, the kernel of $\mathbf{Y}$ is the orthogonal complement of the space spanned by the columns of $\left(\mathbf{g}^{-1}\right)^{\top} \in \mathbb{F}^{nk \times k}$. Thus, $\mathbb{F}^n$ is the direct sum of the kernel of $\mathbf{Y}$ and

the eigenspace of $\mathbf{Y}$ associated with n. Overall, in the absence of noise, the eigenspaces structure of the measurement matrix $\mathbf{Y}$ contains information on $\mathbf{g}$, and therefore contains information on the absolute group elements.

In the presence of noise, the measurement matrix $\mathbf{Y}$ can be viewed as a disrupted form of a hypothetical, clean measurement matrix of the form we described in the previous paragraph. The eigenspaces of the measurement matrix can also be viewed as disruptions of the eigenspaces of this hypothetical matrix. If the disruption is mild enough, the eigenspaces of the top k eigenvalues of $\mathbf{Y}$ is expected to preserve enough of the information of the eigenspace associated with the eigenvalue n of this hypothetical matrix. By projecting these top eigenvectors of the measurement matrix onto the representation of G, it is possible to obtain an estimate of the absolute group elements. This step may also somewhat alleviate the disruption of the eigenspace caused by the noise. The resulting estimation method can be formulated as a two-step procedure:

1. Find a set of unit eigenvectors of $\mathbf{Y}$ associated with its largest k eigenvectors. These can be placed in a matrix $\mathbf{U} \in \mathbb{R}^{nk \times k}$.

2. Project these eigenvectors onto G.

Implementing this two-step method is fairly straightforward using existing numerical algebra libraries. The first step can be easily implemented using all modern numerical eigenproblem solvers. The second step, often referred to as '**rounding**' in the synchronization literature (19; 22), can be more involved, depending largely on G and its chosen representation. A few examples will serve to illustrate the range in the complexity of this task.

EXAMPLE 1. If $G = \mathrm{SO}(2)$ is represented as unit complex numbers, the result of the first step is a unit vector of complex numbers. Simply dividing each of its entries by its complex norm yields a vector of unit complex numbers. This is possible, provided the disruption caused by the noise is mild enough so as to ensure all entries of the eigenvector are non-zero. This rounding procedure was used to great effect in SO(2) synchronization, e.g. (16; 22).

EXAMPLE 2. If $G = \mathrm{SO}(k)$ is represented as real orthogonal $k \times k$ matrices, the result of the first step is a $nk \times k$ matrix which can be treated as a column block matrix with $k \times k$ blocks. The projection is carried out block-wise. The SVD of each block is calculated, yielding a decomposition $\mathbf{B} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^{\top}$, where $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices and $\boldsymbol{\Sigma}$ is a diagonal matrix. The rounding step is completed by replacing every block with the image of the map $\mathbf{B} \mapsto \mathbf{U}\mathrm{diag}\left(1, \ldots, 1, \det\left(\mathbf{U}\mathbf{V}^{\top}\right)\right)\mathbf{V}^{\top}$.

EXAMPLE 3. In (1), $G = \mathrm{SE}(3)$ was represented on the real 4×4 matrices of the form

$$\begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^{\top} & 1 \end{bmatrix}. \tag{3.1}$$

This matrix represents $(R, \mathbf{t}) \in \mathrm{SE}(3)$, where R is a rotation represented as $\mathbf{R}$, a real 3×3 orthogonal matrix with determinant of 1. Consider the matrix $\mathbf{V} \in \mathbb{R}^{4n \times 4}$, the columns of which are the top eigenvectors of the measurement matrix which were calculated in the first step. The rounding procedure considered by (1) begins by extracting every fourth row of $\mathbf{V}$. Denoting the resulting $\mathbb{R}^{n \times 4}$ matrix by $\mathbf{V}_4$, all solutions of $\mathbf{V}_4 \mathbf{u} = \mathbf{0}$ are found by finding an orthonormal basis $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ for the kernel of $\mathbf{V}_4$. A unit vector solution for $\mathbf{V}_4 \mathbf{u}_4 = (1, \ldots, 1)^{\top}$ is also found. Since nothing guarantees that solutions to these equations even exist, they are approximated in the least squares sense. Let $\mathbf{U} \in \mathbb{R}^{4 \times 4}$ be the matrix

with columns $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \mathbf{u}_4$. Let $\mathbf{V}' = \mathbf{VU}$ and substitute every forth row of the resulting matrix with $(0, 0, 0, 1)$. Now, each of the 4×4 blocks of $\mathbf{V}'$ has almost the form (3.1). If this matrix is complex, zero out the imaginary part of every entry. Finally, the top 3×3 submatrix of every block is projected onto SO(3) by utilizing its SVD in the manner described in Example 2.

Non-linear transformations are employed in all three examples. However, the simplicity of the rounding step in Example 1 for SO(2) is contrasted by the multi-step, complicated procedure shown in the Example 3 for SE(3), the focus of this paper. This difference stems from inherent features of the representation of the group. The group SO(2) is embedded within the complex plane in a way which respects the algebraic properties of the representation space. First, the norm of a complex number is directly related to the conjugation operation, i.e. $|z| = \sqrt{zz^*}$ for any $z \in \mathbb{C}$. The submanifold of the complex plane into which SO(2) is embedded, the unit circle, is entirely defined in terms of this conjugation operation, i.e. $\mathbb{T} = \{z \in \mathbb{C} \mid zz^* = 1\}$. Second, the embedding takes into consideration the ∗-algebra structure of the complex plane. The group inverse is the conjugation operation.

These features do not hold for the matrix representation (3.1) of SE(3). As we discuss in greater detail in Appendix A, this representation merely embeds SE(3) within the multiplicative group of the matrix algebra of real 4×4 matrices. It does not rely on or respect the other algebraic structures of the representation space. Merely transposing (3.1) almost always yields a matrix which does not represent an element of SE(3).

We now discuss the nature of the rounding procedure. On the surface, all it does is ensure that the output of the spectral method are elements of the group, but a closer look reveals a more complicated situation. In the noiseless case, which is simplest, we saw that the measurement matrix has a single non-zero eigenvalue. We consider now the linear isometries of its eigenspace with respect to the standard Euclidean norm. These isometries form a compact group. Numerical linear algebra tools ensure we shall obtain unit eigenvectors of the measurement matrix. These may be $\frac{1}{n}\mathbf{g}$, where the $\frac{1}{n}$ term ensures these columns are unit eigenvectors. However, more often they are expected to be $\frac{1}{n}\mathbf{g}\mathbf{S}$, where $\mathbf{S}$ is an isometry of the eigenspace of the non-zero eigenvalue of the measurement matrix.

Let us consider the three examples above. In SO(2) synchronization (Example 1), the group of isometries is SO(2) itself, represented by unit complex numbers. Indeed, since the eigenspace is one-dimensional, every isometry has the form $\mathbf{x} \mapsto \mathbf{x}x$ for $x \in \mathbb{C}$. By substituting $\mathbf{x}x$ into the standard Euclidean norm, it is easy to show that it is an isometry if and only if $\|\mathbf{x}x\| = \|x\|\,|x|$, which holds if and only if $|x| = 1$. Therefore, $x \in$ SO(2). It follows that in the noiseless case, the spectral method recovers the absolute group elements up to a right-multiplication by an element of SO(2). And so, in the presence of noise, the rounding step alleviates a disruption in the eigenspace structure and nothing more.

This is not the case in the other two examples. There, the rounding has to perform other tasks, which are unrelated to the disruption of the eigenspaces caused by the noise and will be required even in its absence. Again, it is instructive to consider the noiseless case first. In SO(k) synchronization (Example 2), the non-zero eigenspace is a k-dimensional subspace of $\mathbb{R}^n$ and its isometry group with respect to the standard Euclidean norm on $\mathbb{R}^n$ is O(k), the group of orthogonal transformations of $\mathbb{R}^k$. The group SO(k) is a proper subgroup of O(k) and so the unit eigenvectors in the eigenspace can be separated into two disjoint populations. The first population is obtained by applying SO(k) to the columns of $\frac{1}{n}\mathbf{g}$ from the right. The second can be obtained by applying elements of O(k) \ SO(k). This means that being a solution to the synchronization problem implies being a basis of an eigenspace of the non-zero eigenvalue of the measurement matrix, but the converse does not hold. We refer to this phenomena as the **eigenspace-synchronization gap**. Therefore, in the presence of noise, the rounding procedure performs at least the following two tasks. It alleviates the disruption to the eigenspace structure *and* ensures that the resulting estimate of the absolute group elements is indeed the absolute group elements

up to a multiplication by an arbitrary element of G from the right. The latter can no longer be taken for granted.

These issues are exacerbated considerably in SE(3) synchronization (Example 3), where the group is non-compact. There, even in the absence of noise, the group SE(3) is not even a subgroup of the isometries of the eigenspace. This is the case because the group of isometries is compact, while SE(3) is not. In this case, the eigenspace-synchronization gap is not so much a gap as an abyss. The domain of the rounding procedure is a compact manifold, matrices formed with four unit eigenvectors as columns. Its codomain is a non-compact manifold, block matrices with blocks of the form (3.1). The rounding procedure maps this compact domain to this non-compact codomain. Thus, even in the absence of noise, for SE(3) synchronization it is another large step away from the neat correspondence exhibited in SO(2) synchronization (Example 1) between eigenvectors of the measurement matrix and estimates of the absolute group elements. In the presence of noise, it does a great deal more than handling the eigenspace disruption.

In closing this subsection, we digress to note an important distinction between the way we introduced the spectral method and the way it is often discussed in the literature. In the literature on synchronization, the relationship between the eigenvectors of the noisy measurement matrix and an estimate of the absolute group elements is often arrived at via optimization terminology. For instance, the initial goal of both (1) and (22) is an estimate the absolute group elements obtained by solving a certain least squares problem over $\widehat{\mathbf{g}} \in G^n$. Here, G^n is the set of all column block matrices with blocks representing elements of G. Since the domain of the problem is non-convex, this is a challenging optimization problem to solve in practice. This constraint is therefore changed to all matrices satisfying $\widehat{\mathbf{g}}^* \widehat{\mathbf{g}} = n\mathbf{I}_k$. In the case of (22), which worked on SO(2) represented on the unit complex numbers, this change in the constraint is a relaxation of the constraint of the original least squares problem. That is, the feasible set of the original problem is replaced by a larger set. However, this change is not a relaxation of a constraint in (1), which worked on SE(3) with the matrix representation (3.1). Regardless of that, in both cases, the new least squares problem turns out to be easily solved by finding the eigenvectors of the measurement matrix.

We note this here merely to point out that the optimization perspective provides a useful heuristic for deriving practically useful estimation methods for synchronization problems. However, it does not offer an immediate way to prove when and why these methods work. We here treat the noisy measurement matrix as a disrupted version of a clean, complete measurement matrix, and seek to approximate the eigenspaces of the latter using the eigenspaces of the former. In the optimization perspective, the very same eigenproblem is considered as a relaxed optimization problem, which one hopes approximates the solution of the non-convex problem. This is a difference without much consequence to the governing mathematics of the spectral method. Regardless of the heuristic used to derive the method, considerable effort is required to come up with and prove theoretical guarantees on the performance of the spectral method. Indeed, there is a rich literature of formal characterizations of when and how solutions of the eigenproblem provide good estimates of the absolute orientations, e.g. (14; 15; 19; 22) with many mathematical tools, including random matrix theory and perturbation theory, being brought to bear on these questions.

3.3 *A spectral approach to* SE(3) *synchronization via dual quaternion representation*

In Section 3.2, we discussed the contrast between the rounding step in SO(2) and SE(3) synchronization in Example 1 and Example 3. We identified two types of differences between them. First, their respective representations have different features. For instance, the representation of SO(2) agreed with involutive algebraic structure of the space it was embedded in, whereas the representation of SE(3) did not. Second,

in a sense, the rounding does more in the SE(3) case than in the SO(2) case. We now utilize the material we surveyed in Section 2 to develop a spectral synchronization method for SE(3), which is more similar to the spectral method for SO(2) in both these respects. We note in passing that the quaternions and the dual quaternions were used in the past to address synchronization problems (9; 24), but this is the first time an analog of the spectral method for synchronization was developed using dual quaternions.

All the necessary components of the spectral method are present over the dual quaternions. Elements of SE(3) are represented as unit dual quaternions, as we surveyed in Section 2.4. This representation agrees with both the multiplication operation of the algebra and its involution, since the dual quaternion conjugate of unit dual quaternion is its inverse. The generalized adjacency matrix of the synchronization graph is a matrix of dual quaternions, rather than a block matrix with real or complex entries. The resulting matrix is Hermitian, in the dual quaternion sense. By the spectral theorem for dual quaternion matrices, it has a eigenvector decomposition, as described in Section 2.5.

Now, given a clean, complete set of measurements, the measurement matrix has the form $\mathbf{Y} = \mathbf{g}\mathbf{g}^*$, where $\mathbf{g} = \left(g_1, \ldots, g_n\right)^\top$ is a column vector of unit dual quaternions and $\mathbf{g}^*$ is its conjugate transpose. In the presence of noise, we consider the dual quaternion eigenvector of the measurement matrix corresponding to its top eigenvalue as an approximation of $\mathbf{g}$, the eigenvector of the clean measurement matrix. Here, 'top' is with respect to the order of eigenvalues established in Theorem 10. This eigenvector can be approximated using the power method we surveyed in Section 2.6. The rounding is carried out by applying the map defined in Theorem 9 to each entry of this eigenvector. This rounding procedure stems directly from the close relationship between the algebraic structure and geometry of the dual quaternion algebra and the unit dual quaternions, as we discussed at the end of Section 2.4. If the noise is mild enough, all entries of the top eigenvector are expected to be invertible, in which case the projection amounts to applying the map $q \mapsto \frac{q}{|q|}$ entry-wise to the eigenvector.

Overall, we obtain the following two-step algorithm:

1. Find the eigenvector of $\mathbf{Y}$ corresponding to its largest eigenvalue, in the order established in Theorem 11. Use the power iteration we surveyed in Section 2.6. Denote this eigenvector by $\widehat{g} = \left(\widehat{g}_1, \ldots, \widehat{g}_n\right)^\top$.

2. Carry out the substitution $\widehat{g}_j \to \frac{\widehat{g}_j}{|g_j|}$ on each of its entries ($j = 1, \ldots, n$).

As we demonstrate in Section 4, this algorithm yields comparable performance with the state-of-the-art spectral approach to SE(3) synchronization described in (1). We emphasize the simplicity of the rounding step compared with the one employed by (1), which we described in Example 3. Its simplicity is owed entirely to the close relationship between SE(3) and the algebra of dual quaternions. We also emphasize that here the rounding procedure does nothing more than handle the disruption the noise causes to the eigenspaces of the clean measurement matrix. Indeed, as we discussed in Section 2.7, the isometry group of the eigenspaces of $\mathbf{Y}$ is actually the unit dual quaternions and as we discussed in Section 2.4, each unit dual quaternion represents an element of SE(3). Thus, there is no eigenspace-synchronization gap, despite the non-compact setting. Our method enjoys similar performance to the current state of the art in spectral synchronization for SE(3), yet reaps the benefits of embedding SE(3) in an algebra that is much closer to it in structure.

Taken together, the resulting spectral method is remarkably similar to the spectral method of SO(2) we showcased in Example 1. This is evident in two places. First, in the properties of the representing space and its relationship to the subset which represents the group. Second, in the underlying tasks performed by the rounding step. This similarity is especially remarkable when we consider the fact

that SO(2) is a compact, commutative group, while SE(3) is a non-compact and non-commutative group.

Finally, we note that (7) addressed a synchronization problem with selection noise, though they do not use the word synchronization at all. They proposed an iterative method, which is distinct from the spectral method we constructed here, both in its form and its theoretical underpinnings. First, it is not the spectral method as we conceived it here. At every iteration, the iteration (2.8) is used to approximate the eigenvector of some matrix of dual quaternions. Second, their iterative method was motivated by the need to solve an explicit optimization problem over the dual quaternions, whereas we relied on the heuristic picture of the eigenspaces of a matrix being disrupted by noise.

4. Numerical experiments

We run several numerical experiments demonstrating the utility of the dual quaternions in synchronization problems. We begin in Section 4.1 by describing how the synthetic synchronization measurement matrices were simulated in our experiments and describe our experimental procedure. In Section 4.2 and Section 4.3 we show and discuss the results of the synthetic data experiments. They demonstrate that using the dual quaternions to represent SE(3) yields comparable estimates with the current state-of-the-art spectral synchronization method (1) while enjoying the simpler rounding procedure obtained by embedding SE(3) in the algebra of dual quaternions. Finally, in Section 4.4, we consider the computational efficiency of our method. We discuss the computational efficiency of our implementation of the algebra of dual quaternions and of the current state of the art. We show that our approach has better computational performance overall than the current state of the art, and that our rounding procedure is faster than the current state of the art by an order of magnitude. We obtained this result even though our implementation is not optimal, which indicates that our approach has untapped potential when it comes to its computational efficiency.

4.1 *Synthetic experiments set-up*

We conducted several experiments on synthetic measurements matrices, modelled after the synthetic data experiments carried out in (1), the current state of the art in SE(3) synchronization. Probability functions on SE(3) were defined using the angle-axis-translation representation discussed in Appendix A. Ground truth elements $g_1, \ldots, g_n \in$ SE(3) were drawn from i.i.d. random variables. The angles were independently and uniformly distributed over $[0, 2\pi)$. The axis was independently and uniformly distributed over S^2. Translations of the form $\mathbf{t} = (t_1, t_2, t_3)^\top \in \mathbb{R}^3$ were independently distributed with t_1, t_2 and t_3 being i.i.d. standard Gaussian random variables. The measurement matrix $\mathbf{Y} \in \mathbb{DH}^{n \times n}$ follows the following model:

$$\mathbf{Y}_{ij} = \begin{cases} e_{ij} \left(s_{ij} g_i g_j^{-1} \varepsilon_{i,j} + \left(1 - s_{ij}\right) c_{ij} \right) & 1 \le i < j \le n, \\ 1 & i = j, \\ \mathbf{Y}_{ji}^* & 1 \le j < i \le n. \end{cases} \tag{4.1}$$

Here, $\left(e_{ij} \mid 1 \le i < j \le n\right)$ are i.i.d Bernoulli random variables with parameter p and $\left(s_{ij} \mid 1 \le i < j \le n\right)$ are i.i.d. are the same with parameter q. The former are indicators of existing measurements and the latter are indicators of entries which were not corrupted. Denoting by $\circ$ the entry-wise product of two matrices of elements of SE(3), the matrices $\mathbf{E} = \left(e_{ij}\right), \mathbf{S} = \left(s_{ij}\right)$ and $\mathbf{E} \circ \mathbf{S} = \left(e_{ij} s_{ij}\right)$ are all the adjacency

matrices of random graphs following the Erdoős–Rényi model with parameter p, q and pq, respectively. The corrupted entries $\left(c_{ij} \mid 1 \leq i < j \leq n\right)$ are i.i.d. with an angle in degrees normally distributed with zero mean and variance σ_r^2, axis distributed uniformly over S^2 and translation with t_1, t_2, t_3 being i.i.d. Gaussian with zero mean and variance σ_t^2. All random variables mentioned above were also independent of each other.

Given a ground truth $\mathbf{g} \in SE(3)^n$ and an estimate $\widehat{\mathbf{g}} \in SE(3)^n$, we measure the quality of the estimate as follows. Due to the symmetry in the observations, we have to align $\mathbf{g}$ and $\widehat{\mathbf{g}}$ in the following sense. We find a $g \in SE(3)$ such that a distance metric between $\mathbf{g}$ and $\widehat{\mathbf{g}}g$ is minimized, where $\widehat{\mathbf{g}}g$ is the result of multiplying by g from the right every entry of $\widehat{\mathbf{g}}$. In Appendix B, we discuss the specifics, including the distance metric we used. Here, we describe the procedure itself. First, we assume that every element of SE(3) is represented by $(q, \mathbf{t})$ with $q \in \mathbb{H}$ being a unit quaternion with non-negative first coordinate as discussed in Section 2.3 and $\mathbf{t} \in \mathbb{R}^3$. Second, if $g = (q, \mathbf{t})$, $\mathbf{g} = \left((q_1, \mathbf{t}_1), \ldots, (q_n, \mathbf{t}_n)\right)$ and $\widehat{\mathbf{g}} = \left((\widehat{q}_1, \widehat{\mathbf{t}}_1), \ldots, (\widehat{q}_n, \widehat{\mathbf{t}}_n)\right)$, we have

$$q = \frac{s}{|s|}, \qquad s = \sum_{j=1}^{n} \widehat{q}_j^* q_j, \tag{4.2}$$

and

$$\mathbf{t} = \frac{1}{n} \sum_{j=1}^{n} \varphi\left(\widehat{q}_j^*\right)\left(\mathbf{t}_j - \widehat{\mathbf{t}}_j\right). \tag{4.3}$$

This alignment procedure was used on the estimate obtained by both methods employed in our experiments, our dual quaternion spectral synchronization and the spectral method of (1), which used the matrix representation (3.1). We note that our alignment method is considerably simpler to apply than the alignment method utilized by (1), whose numerical experiments inspired our own.

All experiments were repeated 50 times, each with freshly generated ground truth and noise realization with the specified experimental parameters. In all our experiments, we used exactly 20 iterations of the dual quaternion power iteration. From each estimate we extracted the entry-wise error, after alignment, for the rotations and translations, separately. For the rotations, the error was calculated as $d\left(q_1, q_2\right) = 2 \arccos\left(2\langle q_1, q_2 \rangle^2 - 1\right)$, where $\langle q_1, q_2 \rangle$ is the Euclidean inner product applied to q_1 and q_2 considered as elements of $\mathbb{R}^4$ with the obvious embedding. See Remark 1 at the end of Appendix B for an explanation of this metric. For the translations, we used the Euclidean distance, $d\left(\mathbf{t}_1, \mathbf{t}_2\right) = \left\|\mathbf{t}_1 - \mathbf{t}_2\right\|$. For each, the mean, minimum and maximum error were calculated and each of these measures was averaged over the 50 repeats. The plotted data for every quantity are always its mean over the repeats.

We described the approach of (1) in Example 3. Both estimation methods were implemented in Python. The original implementation of (1) was in MATLAB. We took care to use equivalent functions to theirs in the Python ecosystem. Code used to simulate the figures below is available at https://github.com/idohadi/dqsync-python/.

4.2 *Response to missing entries and perturbative noise*

In Fig. 2(a), we show the results with increasing levels of multiplicative noise on the rotations and no noise on the translations. In Fig. 2(b), we show the results with increasing levels of multiplicative noise on the translations and no noise on the rotations. Consistent with the semidirect product structure of SE(3),

(a) $\sigma_r = 1°, 2°, \ldots, 20°$ and $\sigma_t = 0$.

(b) $\sigma_r = 0°$ and $\sigma_t = 0.01, 0.02, \ldots, 0.2$.

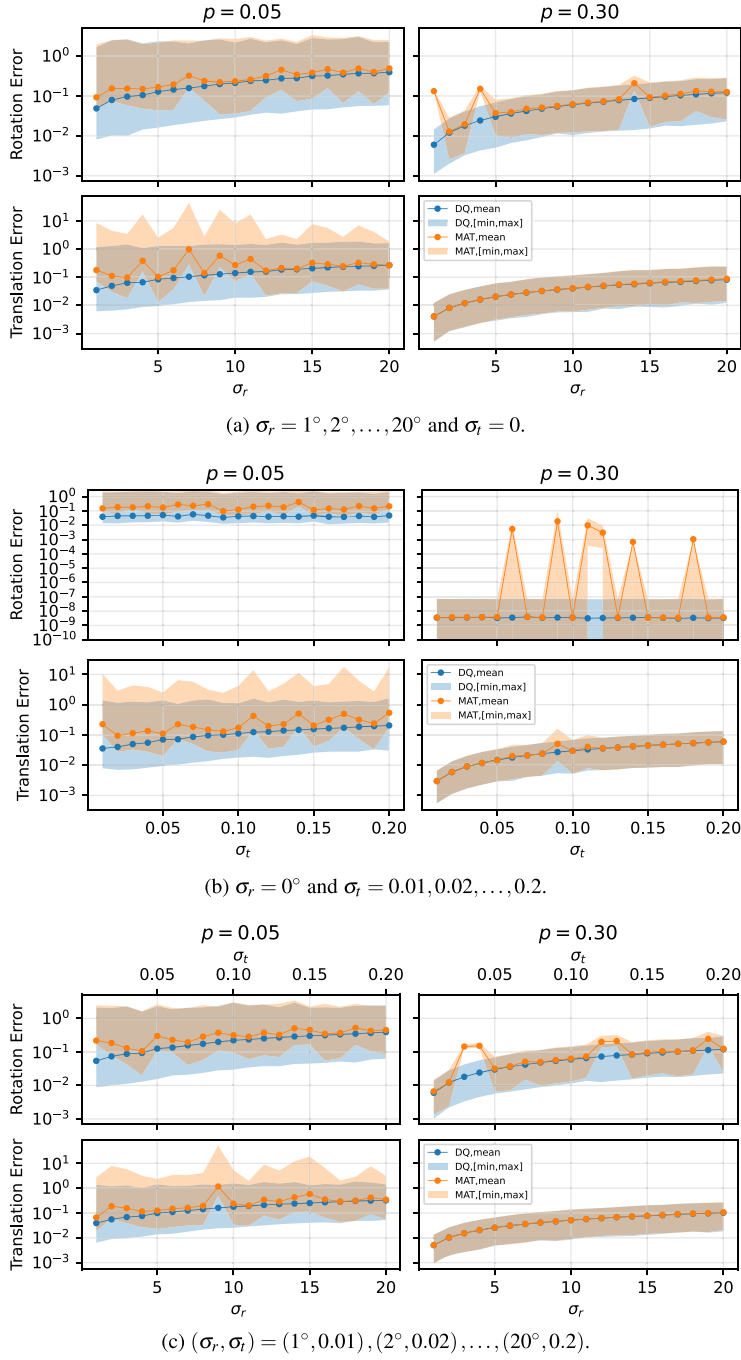(c) $(\sigma_r, \sigma_t) = (1°, 0.01), (2°, 0.02), \ldots, (20°, 0.2)$.

FIG. 2. **Synthetic data experiments with missing entries and perturbative noise.** Throughout experiments, problem size is $n = 100$ and there are no corrupted entries ($q = 1$). Every data point is an average over 50 repeats. The results using the dual quaternion representation are denoted **DQ** and using the matrix representation (3.1) are denoted **MAT**.

in Fig. 2(a) the estimate of the translational component is also affected by noise, whereas in Fig. 2(b) the quality of the estimate of the rotations is basically flat, unaffected by the noise. When we vary the level of noise on both the rotational and translational part in Fig. 2(c), the error rises on both components. Overall, using the dual quaternion representation of SE(3) yields marginally better mean entry-wise error, especially in the estimate of the rotation and when more entries are missing ($p = 0.05$ vs. $p = 0.3$). It is also evident that using the dual quaternions yields somewhat stabler estimates, as indicated by the smoother dual quaternion curves evident in Fig. 2. This may be a feature of our implementation, rather than the method, as these peaks are not evident in the plots shown in (1). All trends we discussed above are also evident in minimum and maximum entry-wise errors. We conclude that the two methods perform comparably and that the dual quaternions can be used effectively to represent SE(3) in applied estimation problems.

4.3   *Response to corruptive noise*

In Fig. 3(a), we used no perturbative noise but there were missing entries. In Fig. 3(b), there are no missing entries but there is perturbative noise. The same trends are evident in both of them. First, the rotation estimate is better using the matrix representation (3.1), while the translation estimate is better using the dual quaternion representation. Second, the differences in estimate quality are larger in Fig. 3(a) than in Fig. 3(b). Overall, these experiments indicate no clear advantage to either method, but do underscore the fact the dual quaternions can be used effectively to represent SE(3) in applied estimation problems.

4.4   *Runtime*

In Fig. 4 we show the runtime of our spectral approach to SE(3) synchronization and that of (1) for increasing problem size *n*. We show the runtime of the eigenproblem solver, the rounding procedure and the total running time which combines both. The eigenproblem solver in our method is the dual quaternion power iteration and an optimized numerical eigenproblem solver provided by the scipy Python library. Each data point is the average running time of 50 repeats. Here, a single repeat starts by generating a fresh synchronization matrix with new ground truth and noise realization with no corruptive noise, $p = 0.05$, $\sigma_r = 20°$ and $\sigma_t = 0.2$. A discussion of our implementation of the algebra of dual quaternions is necessary to fully appreciate these results.

The algebra of real 4×4 matrices, into which (1) embed SE(3), can be implemented using standard numerical algebra libraries and the same holds for block matrices with blocks in this algebra. A single element of this algebra is stored in 16 real floating points numbers. It takes 16 floating point operations to add two elements of this algebra and 112 (4^3 multiplications and $4^2 \cdot 3$ additions) operations to multiply them. In an optimized implementation, only eight real floating points numbers are stored to represent a single dual quaternion. It takes eight floating point operations to add two dual quaternions. Multiplying a pair of quaternions takes 28 (4^2 multiplications and $4 \cdot 3$ additions or subtractions) floating point operations. Therefore, it takes 88 floating points operations to multiply a pair of dual quaternions. Thus, in an optimized implementation, dual quaternions require half the storage, half the floating points operations for addition and roughly a fifth less floating points operations for multiplication, compared with real 4×4 matrices.

In implementing the algebra of dual quaternions, we used the fact that the algebra of dual quaternions is isomorphic to the tensor product $\mathbb{H} \otimes_{\mathbb{R}} \mathbb{D}$ of the algebra of quaternions and the algebra of dual numbers, both algebras over the reals. The algebra of quaternions has a known embedding into the algebra of complex 2×2 matrices (8, Sec. 2.3). The algebra of dual numbers can be realized as the subalgebra of the real 2×2 matrices generated by the identity matrix and a nilpotent 2×2 (satisfying $X^2 = 0$).
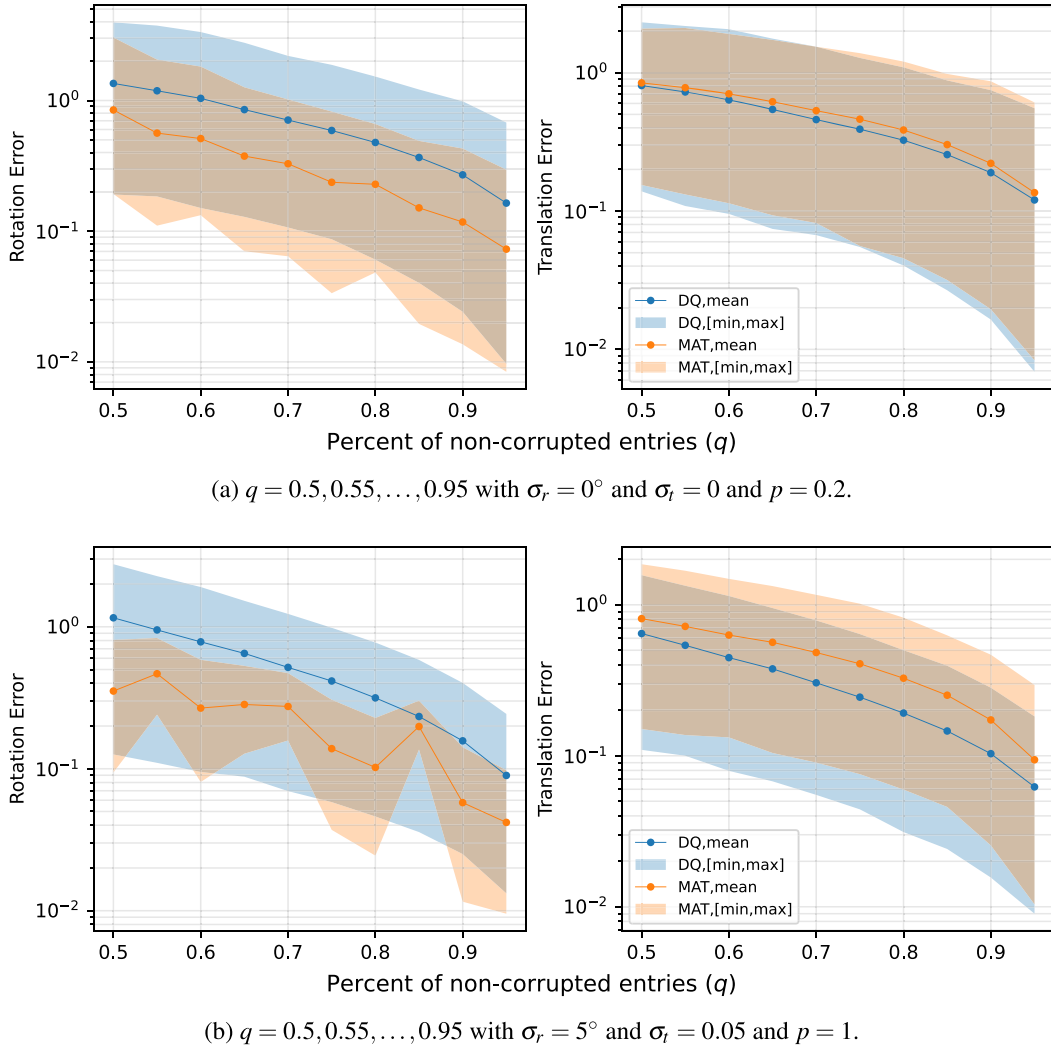
(a) $q = 0.5, 0.55, \ldots, 0.95$ with $\sigma_r = 0°$ and $\sigma_t = 0$ and $p = 0.2$.

(b) $q = 0.5, 0.55, \ldots, 0.95$ with $\sigma_r = 5°$ and $\sigma_t = 0.05$ and $p = 1$.

FIG. 3. **Synthetic data experiments with corruptive noise.** Throughout experiments, problem size is $n = 100$. Every data point is an average over 50 repeats. The results using the dual quaternion representation are denoted **DQ** and using the matrix representation (3.1) are denoted **MAT**.

By choosing explicit generators for these two matrix algebras, the algebra of dual quaternions can then be represented by the algebra generated by the tensor product (that is, the Kronecker product) of these generators. The result is a representation of the dual quaternions on a subalgebra of the complex 4×4 matrices.

The pertinent advantage of this representation of the dual quaternions is the ease of implementation it provides. The matrix addition and product of the complex 4×4 matrices correspond to their dual quaternion counterparts. Thus, using a standard numerical linear algebra library, we can implement algorithms which require addition and multiplication of dual quaternions. Furthermore, modules of

TABLE 1 *Storage and floating point operations for basic operations of various representing spaces of dual quaternions. Under storage we list the number of real numbers stored and floating point operations are the number of operations on real numbers.*

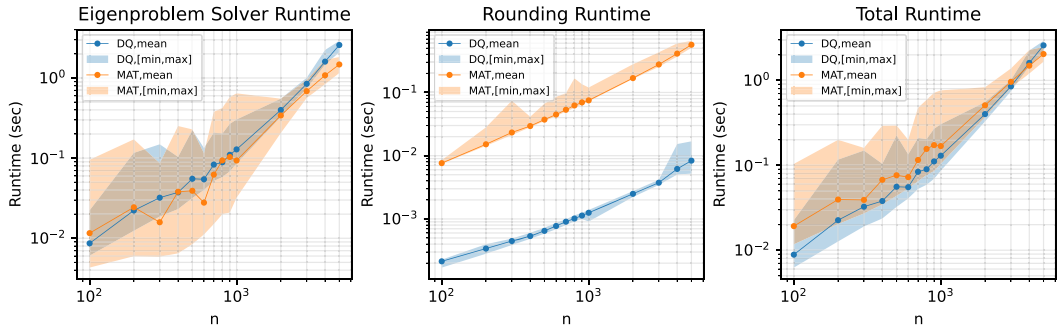| Representing algebra	Storage	Addition	Multiplication
$\mathbb{R}^{4\times4}$	16	16	112
$\mathbb{DH}$ (optimized)	8	8	88
$\mathbb{DH} \subseteq \mathbb{C}^{4\times4}$	32	32	288
botrule			

FIG. 4. Empirical runtime in seconds for the numerical eigenproblem solver (left), rounding procedure (middle) and entire spectral synchronization problem (right) when using the dual quaternion representation (**DQ**) and the matrix representation (3.1) (**MAT**).

vectors or matrices of dual quaternions can also be easily implemented using block matrices with blocks taken from the algebra of complex 4×4 matrices. Its disadvantages are as follows. First, this is an algebra embedding that is not a $*$-algebra embedding. In practice, this means we have to implement a dedicated function for the dual quaternion conjugation operation. Second, and this is most pertinent to our discussion here, the resulting performance is not optimal from a computational efficiency standpoint. A single dual quaternion element is stored as 32 real floating point numbers. Addition requires 32 floating point operations on real numbers. Multiplication requires 288 floating point operations on real numbers. This is considerably worse even than the amount of storage and floating point operations required for basic operations in the matrix embedding of SE(3). In Table 1, we summarize the discussion on the storage and the floating point operations required for various representing spaces of SE(3).

We sought to demonstrate the viability of the dual quaternion embedding for spectral synchronization, rather than its superior computational performance. Our implementation of the algebra of dual quaternions was therefore unoptimized. Yet, viewing the results in Fig. 4 with that knowledge in hand, we see that even our unoptimized implementation of the dual quaternions challenges the current state of the art in spectral synchronization for SE(3). The rounding procedure is consistently an order of magnitude faster. In addition, the dual quaternion power iteration has inferior performance to the optimized eigenproblem solver used in the matrix representation, but the two are fairly close. Owing mainly to its superior performance in the rounding step, the total running time of the dual quaternion approach to synchronization is actually better. Only when the problem size is around $n = 5000$, the matrix embedding provides better performance, due to the better performance of its eigenproblem solver.

Beyond the superior performance demonstrated already, taking these results in context suggests that a dual quaternion approach to spectral synchronization has the potential to be even more computationally efficient than the current state of the art, including for larger problem sizes.

Data availability

The data underlying this article are available in the GitHub repository `dqsync-python` at https://github.com/idohadi/dqsync-python/.

Funding

1. ARRIGONI, F., ROSSI, B. & FUSIELLO, A. (2016) Spectral synchronization of multiple views in se(3). *SIAM J. Imaging Sci.*, **9**, 1963–1990.
2. ARRIGONI, F., ROSSI, B., FRAGNETO, P. & FUSIELLO, A. (2018) Robust synchronization in so (3) and se (3) via low-rank and sparse matrix decomposition. *Comput. Vision Image Understanding*, **174**, 95–113.
3. BENDORY, T., BARTESAGHI, A. & SINGER, A. (2020) Single-particle cryo-electron microscopy: mathematical theory, computational challenges, and opportunities. *IEEE Signal Process. Mag.*, **37**, 58–76.
4. BLACKETT, D. W. (1956) The near-ring of affine transformations. *Proc. Am. Math. Soc.*, **7**, 517–519.
5. BOUMAL, N. (2016) Nonconvex phase synchronization. *SIAM J. Optim.*, **26**, 2355–2377.
6. BOUMAL, N., AMIT SINGER, P.-A. & ABSIL, & VINCENT D. BLONDEL. (2014) Cramrrao bounds for synchronization of rotations. *Inf. Inference*, **3**, 1–39.
7. CUI, C. & QI, L. (2024) A power method for computing the dominant eigenvalue of a dual quaternion hermitian matrix. *J. Sci. Comput.*, **100**, 21.
8. GARLING, D. J. H. (2011) *Clifford Algebras: An Introduction*, vol. **78**. Cambridge University Press, New York, USA.
9. GOVINDU, V.M. (2001) Combining two-view constraints for motion estimation. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001,* volume **2**, pages II–II.
10. HAZEWINKEL, M., GUBARENI, N. & KIRICHENKO, V. V. (2004) *Algebras, Rings and Modules*, vol. **1**. Springer Science & Business Media.
11. HUYNH, D. Q. (2009) Metrics for 3d rotations: comparison and analysis. *J. Math. Imaging Vision*, **35**, 155–164.
12. JANCO, N. & BENDORY, T. (2023) Unrolled algorithms for group synchronization. *IEEE Open J. Signal Processing*. **4**, 424–436.
13. LERMAN, G. & SHI, Y. (2021) Robust group synchronization via cycle-edge message passing. *Found. Comput. Math.*, 1–77.
14. LING, S. (2022) Improved performance guarantees for orthogonal group synchronization via generalized power method. *SIAM J. Optim.*, **32**, 1018–1048.
15. LING, S. (2022) Near-optimal performance bounds for orthogonal and permutation group synchronization via spectral methods. *Appl. Comput. Harmon. Anal.*, **60**, 20–52.
16. PERRY, A., WEIN, A. S., BANDEIRA, A. S. & MOITRA, A. (2018) Message-passing algorithms for synchronization problems over compact groups. *Comm. Pure Appl. Math.*, **71**, 2275–2322.
17. QI, L. & LUO, Z. (2021) *Eigenvalues and singular values of dual quaternion matrices*.
18. QI, L., LING, C. & YAN, H. (2022) Dual quaternions and dual quaternion vectors. *Commun. Appl. Math. Comput.*, **4**, 1494–1508.
19. ROMANOV, E. & GAVISH, M. (2020) The noise-sensitivity phase transition in spectral group synchronization over compact groups. *Appl. Comput. Harmon. Anal.*, **49**, 935–970.

20. SELIG, J. M. (2005) *Geometric Fundamentals Of Robotics*, vol. **128**. Springer, New York, USA.
21. SHKOLNISKY, Y. & SINGER, A. (2012) Viewing direction estimation in cryo-EM using synchronization. *SIAM J. Imaging Sci.*, **5**, 1088–1110.
22. SINGER, A. (2011) Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.*, **30**, 20–36.
23. THRUN, S., BURGARD, W. & FOX, D. (2005) *Probabilistic Robotics*. London, England August: Intelligent Robotics and Autonomous Agents series. MIT Press.
24. TORSELLO, A., RODOL, E. & ALBARELLI, A. Multiview registration via graph diffusion of dual quaternions. Conference on Computer Vision and Pattern Recognition (CVPR) **2011**, 2441–2448.
25. TREFETHEN, L. N. & BAU, D. (1997) *Numerical Linear Algebra*. SIAM.
26. WANG, L. & SINGER, A. (2013) Exact and stable recovery of rotations for robust synchronization. *Inf. Inference*, **2**, 145–193.
27. ZYEIL, O., VORONINSKI, V., BASRI, R. & SINGER, A. (2017) A survey of structure from motion. *Acta Numer.*, **26**, 305–364.

A. Comparison of representations of SE(3)

We compare the properties of three representations of SE(3) which were used in applied mathematics work. Our discussion is skewed towards properties which we think are central to how notions from linear algebra can be used to solve applied problems over SE(3).

The most straightforward way to represent elements of SE(3) is as affine transformations, a pair $(R, \mathbf{t}) \in \mathrm{SO}(3) \ltimes \mathbb{R}^3$. These may come in several variants, differing in the way the rotation is represented. Here we survey the axis-angle and rotation matrix representations.

In the axis-angle representation, a rotation R is represented by the axis $\mathbf{v} \in S^2 \subset \mathbb{R}^3$ of the rotation and the angle θ (in radians) by which vectors are rotated around $\mathbf{v}$. Both can be concisely encapsulated in a **rotation vector** $\mathbf{x} = \theta \mathbf{v} \in \mathbb{R}^3$, which are elements of $B_3(\mathbf{0}, 2\pi)$, where $B_3(\mathbf{x}, r) = \{\mathbf{y} \in \mathbb{R}^3 \mid \|\mathbf{x} - \mathbf{y}\|_2 < r\}$. Every pair of a rotation vector and a translation can represent an element of SE(3) and one can easily extend that to pairs of the form $\mathbb{R}^3 \times \mathbb{R}^3$ by taking the Euclidean norm of the rotation vector modulo 2π. Thus, we say that **representation space**, the space of all elements of the form of the given representation, is $\mathbb{R}^3 \times \mathbb{R}^3$. The **representing subset**, the subset of the representation space into which SE(3) is embedded, is $B_3(\mathbf{0}, 2\pi) \times \mathbb{R}^3$. Furthermore, there is projection from the representation space to the representation subset, mapping every element of the representation space to the representing subset. We refer to this representation as the **axis-angle-translation representation**.

From all this it follows in particular that in this representation the dimension of the representation space is 6, which matches the manifold dimension of SE(3), which makes it the most parsimonious representation we will survey. However, there is no direct way to calculate the product of two elements of SE(3) in this representation. It is necessary to convert them into one of the other representations we will survey. We do not know of any established way to construct something like an algebra over this representation space and consequently we know of no such way to construct matrix algebras and investigate their spectral properties.

The other SO(3) representation we consider are rotation matrices, real 3×3 orthogonal matrices with a determinant of 1. This representation comes with a multiplicative structure of matrix multiplication and it is embedded in the matrix algebra of all real 3×3 matrices. The representation space of SE(3) is then $\mathbb{R}^{3 \times 3} \times \mathbb{R}^3 \cong \mathbb{R}^{12}$. Its multiplicative and vector space structure can be identified with the set of all affine transformations of $\mathbb{R}^3$, transformations of the form $\mathbf{x} \mapsto A\mathbf{x} + \mathbf{t}$, where A is a linear transformation and $\mathbf{t} \in \mathbb{R}^3$. This set has both an addition of and multiplication operations, defined by adding or composing, respectively, two affine transformations. Despite this, it does not form a ring, because these two binary

operations do not satisfy one of the distributive properties. Borrowing an example from (4), suppose that $(A_1, \mathbf{t}_1)$ and $(\mathbf{0}, \mathbf{t}_2)$ are two affine transformations. Here, $\mathbf{0}$ is the zero linear transformation and $\mathbf{t}_2 \neq 0$. We have

$$(\mathbf{0}, \mathbf{t}_2) \circ \big((A_1, \mathbf{t}_1) + (A_1, \mathbf{t}_1)\big) = (\mathbf{0}, \mathbf{t}_2),$$

but

$$(\mathbf{0}, \mathbf{t}_2) \circ (A_1, \mathbf{t}_1) + (\mathbf{0}, \mathbf{t}_2) \circ (A_1, \mathbf{t}_1) = 2\mathbf{t}_2.$$

While not being a ring, the affine transformations form a different algebraic structure, that of a near-ring. As far as we know, matrix spaces and spectral decompositions were not studied over near-rings in general and over the affine transformations in particular. Furthermore, to the best of our knowledge, a projection of affine transformations onto the representing subset has not been worked out in the literature.

Finally, SE(3) has a group embedding within the matrix algebra $\mathbb{R}^{4 \times 4}$. An element $(R, \mathbf{t}) \in$ SE(3), where R is represented as $\mathbf{R}$, real 3×3 orthogonal matrix with determinant of 1, can be represented as a block matrix

$$\begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix}. \tag{A.1}$$

Rectangular block matrix where all blocks are of this form obviously belong to the matrix algebra $\mathbb{R}^{4n \times 4n}$, which definitely has a spectral decomposition theorem and many other matrix decompositions. Furthermore, an invertible $\mathbf{A} \in \mathbb{R}^{4 \times 4}$ can also be projected onto the representing subset by a two-step procedure. First, finding a invertible matrix $\mathbf{B} \in \mathbb{R}^{4 \times 4}$ such that the matrix $\mathbf{AB}$ has $(0, 0, 0, 1)$ in its bottom row. This amounts to changing the coordinates of the matrix $\mathbf{A}$ to one which satisfies this particular constraint on the bottom row. Second, denoting by $(\mathbf{BA})_{3 \times 3}$ the upper left 3×3 submatrix of $\mathbf{BA}$, one takes its singular value decomposition $(\mathbf{BA})_{3 \times 3} = \mathbf{VSU}^\top$, where $\mathbf{S}$ is diagonal and $\mathbf{U}$ and $\mathbf{V}$ are orthogonal. To obtain a matrix of the form (A.1), replace $(\mathbf{BA})_{3 \times 3}$ with $\mathbf{V}\mathrm{diag}\,(1, 1, \det{(\mathbf{VU})})\,\mathbf{V}^\top$.

The great advantage of the matrix representation (A.1) is the fact that its representation space is a matrix algebra. This allows the use of the familiar linear algebraic toolkit and in particular allows one to define the projection we described above. However, this is also the source of the greatest disadvantage. The group SE(3) is embedded in the multiplicative group of $\mathbb{R}^{4 \times 4}$, and no more. The other algebraic structures defined on $\mathbb{R}^{4 \times 4}$ are inconsistent with the embedding. Merely transposing (A.1) would generally take one out of the representing subspace, indicating the involution of $\mathbb{R}^{4 \times 4}$ plays no role in representing SE(3). Furthermore, the relationship between the projection and group structure is hard to characterize. This last problem comes to the fore in the situation discussed in Section 3.2, where we describe the projection onto SE(3) of an entire column block matrix with $\mathbb{R}^{4 \times 4}$ blocks, which is used in an actual application (1). Perhaps underscoring the gap between SE(3) and the representation space in this case is the fact SE(3) is a manifold of dimension 6, while the representation space is 16-dimensional.

Consider the unit dual quaternion representation of SE(3) we surveyed in Section 2.4. We argue that the representation space is very close to the structure of SE(3), yet similar enough to ordinary vector spaces to allow some of the same linear-algebraic notions that are so useful to the applied mathematician to be brought to bear on applied problems. The representation space is very close to the structure of SE(3) in two ways. First, by the fact nearly all dual quaternions are elements of SE(3) scaled by dual numbers, as shown by Theorem 9 of (7). Second, the inverse of an element of SE(3) represented on the

unit dual quaternion is its dual quaternion conjugate, as we saw Section 2.4. In this way, the involutive algebra structure of the dual quaternions is intimately tied with the group structure of SE(3). Most importantly, despite the non-commutativity of this algebra, it is sufficiently simple to allow one to establish a spectral theorem as we saw in Section 2.5 and also the derivation of a numerical schemes to approximate eigenvectors, like we surveyed in Section 2.6. We suggest that other numerical schemes, like the Arnoldi iteration and other Krylov subspace methods, may also be generalizable to the dual quaternions.

Taken together, these properties indicate that the dual quaternion algebra balances well the difficulties incurred by working over non-commutative algebra with the advantages it offers.

B. Alignment of two dual quaternion vectors

Let $\mathbf{g} = \left(g_1, \ldots, g_n\right)^\top \in \mathrm{SE}(3)^n$ and $\widehat{\mathbf{g}} = \left(\widehat{g}_1, \ldots, \widehat{g}_n\right)^\top \in \mathrm{SE}(3)^n$. Assume that elements of SE(3) are represented as $g_j = \left(q_j, \mathbf{t}_j\right)$ and $\widehat{g}_j = \left(\widehat{q}_j, \widehat{\mathbf{t}}_j\right)$, where q_j and $\widehat{q}_j$ are unit quaternions with a non-negative first coordinate, representing a rotation as we discuss in Section 2.4, and $\mathbf{t}_j$ and $\widehat{\mathbf{t}}_j$ are translations represented as vectors in $\mathbb{R}^3$. Recall proposition 7 and the double cover φ defined there. Aligning $\mathbf{g}$ and $\widehat{\mathbf{g}}$ amounts to finding $g = (q, \mathbf{t})$ solving

$$\min_{g \in \mathrm{SE}(3)} \left\{ \sum_{j=1}^n \left\| \widehat{\mathbf{q}}_j \mathbf{q} - \mathbf{q}_j \right\|_F^2 + \sum_{j=1}^n \left\| \varphi(\widehat{q}_j)(\mathbf{t}) + \widehat{\mathbf{t}}_j - \mathbf{t}_j \right\|_2^2 \right\}. \tag{B.1}$$

Here, $\widehat{\mathbf{q}}_j$, $\mathbf{q}_j$ and $\mathbf{q}$ are real 3×3 orthogonal matrices with determinant of 1 representing the same elements of SO(3) as $\widehat{q}_j$, q_j and q. These matrices are constructed by representing $\varphi(q)$ in the standard basis of $\mathbb{R}^3$. Also, $\|\cdot\|_2$ is the Euclidean norm and $\|\cdot\|_F$ is the Frobenius norm. Solving (B.1) amounts to finding g minimizing the sum of squared distances between corresponding rotations of $\mathbf{g}$ and $\widehat{\mathbf{g}}$ and a sum of squared distances between the translations of $\mathbf{g}$ and $\widehat{\mathbf{g}}$.

The two sums in (B.1) are two decoupled optimization problems, since the left one depends only on q and the right one only on $\mathbf{t}$. We consider these two decoupled problems separately. We begin with the translation sum. Let $F(\mathbf{t})$ be the right sum in (B.1). We wish to find the minimum of F in $\mathbb{R}^3$. Because the Euclidean norm is preserved under orthogonal transformations and φ is a homomorphism, F can be written as

$$F(\mathbf{t}) = \sum_{j=1}^n \left\| \mathbf{t} + \varphi\left(\widehat{q}_j^*\right)\left(\widehat{\mathbf{t}}_j - \mathbf{t}_j\right) \right\|^2.$$

Fermat's theorem then yields that its exterema is

$$\mathbf{t} = \frac{1}{n} \sum_{j=1}^n \varphi\left(\widehat{q}_j^*\right)\left(\mathbf{t}_j - \widehat{\mathbf{t}}_j\right),$$

exactly as in (4.3). That it is a minimum follows from the fact F is convex and defined on $\mathbb{R}^3$, and therefore can have no local maxima.

The rotation sum is slightly more involved. First, we note that for any $\mathbf{R}_1$ and $\mathbf{R}_2$ two real 3×3 orthogonal matrices with determinant of 1 we have

$$\left\| \mathbf{R}_1 - \mathbf{R}_2 \right\|_F^2 = \left\| \mathbf{R}_1 \right\|_F^2 + \left\| \mathbf{R}_2 \right\|_F^2 - 2 \operatorname{tr} \left(\mathbf{R}_1 \mathbf{R}_2^\top \right) = 3 - 2 \operatorname{tr} \left(\mathbf{R}_1^\top \mathbf{R}_2 \right).$$

A minimizer of the left sum in (B.1) is therefore a maximizer over the unit quaternions with non-negative first coordinate of

$$G_1(\mathbf{q}) = \sum_{j=1}^n \operatorname{tr} \left(\mathbf{q}^\top \widehat{\mathbf{q}}_j^\top \mathbf{q}_j \right).$$

Finally, let q_1 and q_2 are quaternions representing the same elements of SO(3) as $\mathbf{R}_1$ and $\mathbf{R}_2$, respectively. As we argue below, the following identity holds:

$$\operatorname{tr} \left(\mathbf{R}_1^\top \mathbf{R}_2 \right) = 4 \langle q_1, q_2 \rangle^2 - 1, \tag{B.2}$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product on $\mathbb{R}^4$ applied to quaternions via their obvious identification with $\mathbb{R}^4$. Therefore, the maximizer of G_1 over the unit quaternions with non-negative first coordinate is the maximizer over the same domain of

$$G(q) = \sum_{j=1}^n \left\langle \widehat{q}_j q, q_j \right\rangle.$$

We note that

$$\langle q_1, q_2 \rangle = \frac{1}{2} \left(q_1^* q_2 + q_2^* q_1 \right),$$

and so

$$\langle q_1 q, q_2 \rangle = \frac{1}{2} \left(q^* q_1^* q_2 + q_2^* q_1 q \right) = \frac{1}{2} \left(q^* \left(q_1^* q_2 \right) + \left(q_1^* q_2 \right)^* q \right) = \langle q, q_1^* q_2 \rangle.$$

Therefore, for any unit quaternion with non-negative first coordinate q we have

$$G(q) = \sum_{j=1}^n \left\langle q, \widehat{q}_j^* q_j \right\rangle = \left\langle q, \sum_{j=1}^n \widehat{q}_j^* q_j \right\rangle \le \|q\|_2 \left\| \sum_{j=1}^n \widehat{q}_j^* q_j \right\|_2 = \left\| \sum_{j=1}^n \widehat{q}_j^* q_j \right\|_2.$$

In the above we used the Cauchy–Schwartz inequality applied to the quaternions, again identified as elements of $\mathbb{R}^4$. Simple substitution into G shows that the q defined in (4.2) is a unit quaternion with non-negative first coordinate which achieves this bound.

It remains to prove (B.2). It follows from the following lemma:

LEMMA 12. Let q be a unit quaternion. Let $\mathbf{R}$ be the real 3×3 orthogonal matrix with determinant of 1 formed by representing $\varphi(q)$ using the standard basis of $\mathbb{R}^3$. If $q = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$, then

$$\operatorname{tr}(\mathbf{R}) = 4a^2 - 1.$$

We obtain (B.2) as a corollary of Lemma 12 due to a combination of two facts. First, $\mathbf{R}_1^\top \mathbf{R}$ in (B.2) expressed in quaternions is $q_1^* q_2$. Second, as can be worked out from (2.4), the first coordinate of the product $q_1^* q_2$ is $\langle q_1, q_2 \rangle$.

Proof. Using the axis-angle representation of $\mathbf{R}$, we consider it a rotation around axis $\mathbf{r} \in S^2$ by angle θ. From (11, Appendix A), $\operatorname{tr}(\mathbf{R}) = 2\cos\theta + 1$, while (11, eq. (6)) yields $a = \cos\frac{\theta}{2}$. Therefore, using the trigonometric identity $\cos\theta = 2\cos^2\frac{\theta}{2} - 1$, we obtain

$$4a^2 - 1 = 4\cos^2\frac{\theta}{2} - 1 = 2\left(2\cos^2\frac{\theta}{2} - 1\right) + 1 = 2\cos\theta + 1,$$

as required. $\square$

REMARK 1. From the identities used in the proof of Lemma 12, it follows that the metric we defined on the unit quaternions in Section 4.1 satisfies $d(q_1, q_2) = 2\arccos(\cos\theta) = 2\theta_{12}$, where $\theta_{12} \in [0, \pi)$ and θ is the rotation angle in the angle-axis representation of the rotation $q_1^* q_2$.