Rank computation in Euclidean Jordan algebras

Michael Orlitzky

April 8, 2022

Abstract

Euclidean Jordan algebras are the abstract foundation for symmetric cone optimization. Every element in a Euclidean Jordan algebra has a complete spectral decomposition analogous to and subsuming that of a real symmetric matrix into rank-one projections. This general spectral decomposition stems from the element's likewise-analogous characteristic polynomial whose degree (they all have the same degree) is called the rank of the algebra. As a prerequisite for the spectral decomposition, we derive an algorithm that computes the rank of a Euclidean Jordan algebra and allows us to construct the characteristic polynomials of its elements. The ultimate goal of this work is to support a generic computational framework for solving symmetric cone optimization problems in Jordan-algebraic terms.

Keywords: Euclidean Jordan algebra, characteristic polynomial, spectral decomposition, rank, symmetric cone

MSC2020: 17–08, 17C99, 13P15, 15A06, 15A54

1 Introduction

Jordan algebras began as a formalism for quantum mechanics that was intended to replace the Copenhagen interpretation [10]. *Formally-real* Jordan algebras, specifically, were modeled after the properties enjoyed by Hermitian matrices, but were soon found to be lacking when Jordan, von Neumann, and Wigner [8] classified the finite-dimensional specimens up to isomorphism. This classification was "deeply disappointing to physicists," since it showed that no finite-dimensional algebra was a suitable home for quantum mechanics [10]. Only many years later would it prove to be a boon for optimizers.

A *Euclidean Jordan algebra* is a finite-dimensional formally-real Jordan algebra additionally endowed with an "associative" inner product [3].

Definition 1. A *Euclidean Jordan algebra* $(V, \circ, \langle \cdot, \cdot \rangle)$ consists of a finite-dimensional real Hilbert space $(V, \langle \cdot, \cdot \rangle)$ and a commutative bilinear *Jordan product* $\circ$ such that

$$\forall x, y \in V : x \circ ((x \circ x) \circ y) = (x \circ x) \circ (x \circ y),$$

$$\forall x, y, z \in V : \langle x \circ y, z \rangle = \langle y, x \circ z \rangle ,$$

and having a multiplicative unit element $1_V \in V$ such that

$$\forall x \in V : 1_V \circ x = x.$$

Euclidean Jordan algebras arise in optimization because many of the cones used in conic optimization are the cone of squares in some Euclidean Jordan algebra. After publication in 1984 [9], Karmarkar's interior point methods were quickly extended to second-order and semidefinite cone programs—special cases that can be solved efficiently. Nesterov and Nemirovskii [11] showed in 1994 that a general cone program can be solved efficiently if one knows a *self-concordant* barrier function for the cone, and this coincided fortuitously with the publication of Faraut and Korányi's *Analysis on Symmetric Cones*, likely the most comprehensive and oft-cited reference on Euclidean Jordan algebras [3]. Shortly thereafter, Güler [7] pointed out that self-concordant barrier functions are known for the self-dual homogeneous (now called *symmetric*) cones that are nothing other than the cones of squares in Euclidean Jordan algebras.

Since Jordan algebras were modeled upon the Hermitian matrices, it is no surprise that the operation $X \circ Y \coloneqq (XY + YX)/2$ and the trace inner-product turn the spaces of n-by-n real-symmetric or complex-Hermitian matrices into Euclidean Jordan algebras whose unit element is the identity matrix and whose cone of squares is the positive-semidefinite cone. Particularly for us, these will serve as canonical examples of Euclidean Jordan algebras. Here is another.

Example 1 (Jordan spin algebra)**.** In $V = \mathbb{R}^n$ with the usual inner product, let $x \coloneqq (x_1, \bar{x})^T \in \mathbb{R} \times \mathbb{R}^{n-1}$ be written in block form and likewise for y. Then

$$x \circ y \coloneqq \begin{bmatrix} x_1 \\ \bar{x} \end{bmatrix} \circ \begin{bmatrix} y_1 \\ \bar{y} \end{bmatrix} = \begin{bmatrix} \langle x, y \rangle \\ y_1 \bar{x} + x_1 \bar{y} \end{bmatrix}$$

is a commutative bilinear operation with unit element $1_V = \left(1, \bar{0}\right)^T$ satisfying Definition 1. As a result, $(\mathbb{R}^n, \circ, \langle \cdot, \cdot \rangle)$ forms a Euclidean Jordan algebra known as the *Jordan spin algebra*. Its cone of squares is the second-order cone.

Motivated by the connection between symmetric cones and Euclidean Jordan algebras, Faybusovich embarked on a quest to show that certain interior-point methods can be described directly in Jordan-algebraic terms [4, 5, 6]. As a result, we can now solve a broad class of optimization problems by performing computations in Euclidean Jordan algebras. Many of these computations are tractable because the algebra is finite-dimensional and the Jordan product is bilinear. The spectral decomposition, in contrast, remains slippery.

Theorem 1 (Faraut and Korányi [3], Theorems III.1.1–2)**.** *If $(V, \circ, \langle \cdot, \cdot \rangle)$ is a Euclidean Jordan algebra of rank r and if $x \in V$, then there exists a Jordan frame $\{c_1, c_2, \ldots, c_r\}$ in V and real numbers $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_r$ such that*

$$x = \lambda_1 c_1 + \lambda_2 c_2 + \cdots + \lambda_r c_r.$$

In this *spectral decomposition* of x, the λ_i are called the *eigenvalues* of x and they are independent of the Jordan frame used for the decomposition. (The meaning of "Jordan frame" will not be important.) The reader will notice the similarity of this result to the usual spectral decomposition of a Hermitian matrix into rank-one spectral projectors. Our decomposition is unique in the same precise technical sense—after grouping eigenvalues, ignoring rearrangements, and so on.

The rank r of the algebra that appears in Theorem 1 deserves an explanation. In any Euclidean Jordan algebra, we can take powers of a single element x to obtain $x^2 \coloneqq x \circ x$ and so forth. We define x^0 to be the unit element of the algebra for consistency. After a certain point, the set of these powers must become linearly-dependent, because the ambient vector space is finite-dimensional. Thus, eventually, for some $d \in \mathbb{N}$, we can write x^d as a linear combination of $\left\{x^0, x^1, \ldots, x^{d-1}\right\}$. The *minimal polynomial* of x is the unique monic polynomial of minimal degree that evaluates to zero on x. In other words, the minimal polynomial of x shows how you would write x^d as a linear combination of lower powers, where d is the smallest natural number making it possible to do so. In the algebras of real or complex Hermitian matrices, this definition coincides with the usual linear-algebraic definition of minimal polynomial.

Definition 2. If $(V, \circ, \langle\cdot,\cdot\rangle)$ is a Euclidean Jordan algebra and if $x \in V$, then the *degree of* x is the degree of the minimal polynomial of x, and is written $\deg(x)$. The *rank* of the algebra V is the maximal degree of its elements,

$$\operatorname{rank}(V) \coloneqq \max\left(\{\deg(x) \mid x \in V\}\right).$$

If $\deg(x) = \operatorname{rank}(V)$, we say that x is a *regular element* of the algebra.

The rank is well-defined because it's a natural number bounded above by the dimension of the ambient vector space. But this definition doesn't tell us how we might *find* either the rank or a regular element. This is annoying because the rank is a prerequisite for the spectral decomposition in Theorem 1.

In many cases, the rank is known. In the algebras of n-by-n real or complex Hermitian matrices, Theorem 1 reduces to the usual spectral decomposition and the rank of the algebra is the number of rank-one projections you get, namely n. The rank of the Jordan spin algebra is also easy to compute.

Example 2 (spin algebra rank)**.** If $x = (x_1, \bar{x})^T$ is an arbitrary element of the Jordan spin algebra from Example 1, then

$$x^0 = \begin{bmatrix} 1 \\ \bar{0} \end{bmatrix}, \quad x^1 = \begin{bmatrix} x_1 \\ \bar{x} \end{bmatrix}, \quad x^2 = \begin{bmatrix} \|x\|^2 \\ 2x_1\bar{x} \end{bmatrix},$$

and x^2 is a linear combination of x^0 and x^1:

$$x^2 = \left(\|x\|^2 - 2x_1^2\right) x^0 + 2x_1 x^1.$$

As a result, the minimal polynomial of x has degree two at most. We may need fewer powers for some x, but we cannot need more. Now let $e = (1, 1, \ldots, 1)^T$

be the vector of ones. For this element,

$$e^0 = \begin{bmatrix} 1 \\ \bar{0} \end{bmatrix} \text{ and } e^1 = \begin{bmatrix} 1 \\ \bar{e} \end{bmatrix}.$$

Clearly, e^1 cannot be written as a linear combination (that is, a scalar multiple) of e^0, nor vice-versa. So, $\deg(e) > 1$. We showed that the maximum degree of any element in this algebra is two, and then we found an element of degree two or more. It follows that the rank of the Jordan spin algebra is two.

In fact, the ranks of all "simple" Euclidean Jordan algebras are known, and any particular algebra is isomorphic to an orthogonal direct sum of simple factors [3]. In theory this determines the rank of any algebra, but only if you know the isomorphism that allows you to identify the factors.

As in the motivating case of the Hermitian matrices, the spectral decomposition in a Euclidean Jordan algebra is closely related to the *characteristic polynomial* of an element. Taking a cue from standard linear algebra, the characteristic polynomial of a regular element in a Euclidean Jordan algebra is defined to be equal to its minimal polynomial. A "characteristic polynomial of" function can then be defined on the entire algebra by continuously extending the "minimal polynomial of" function from the dense set of regular elements to the entire algebra. The eigenvalues of the element in Theorem 1 turn out to be the roots of its characteristic polynomial. Below we paraphrase the relevant portions of Faraut and Korányi's Proposition II.2.1 [3].

Theorem 2. *If $(V, \circ, \langle\cdot,\cdot\rangle)$ is a Euclidean Jordan algebra of rank r and dimension n, then there exists a regular element $\xi \in V$, a basis $\mathbf{b}_\xi$, and polynomials a_0 through a_{r-1} in $\mathbb{R}[X_1, X_2, \ldots, X_n]$ such that the characteristic polynomial of any $x \in V$ is*

$$\Lambda^r + \sum_{i=0}^{r-1} \widehat{a_i}\left(\mathbf{b}_\xi(x)\right)\Lambda^i \in \mathbb{R}[\Lambda].$$

Here we have used $\mathbf{b}_\xi(x)$ to denote the representation of the algebra element x with respect to the basis $\mathbf{b}_\xi$, and $\widehat{a_i} : \mathbb{R}^n \to \mathbb{R}$ to denote the polynomial function that corresponds to $a_i \in \mathbb{R}[X_1, X_2, \ldots, X_n]$.

The basis $\mathbf{b}_\xi$ in this result begins with the first r powers $\{\xi^0, \xi^1, \ldots, \xi^{r-1}\}$ of the regular element ξ, guaranteed to be linearly-independent by the definitions of degree and rank. This choice expedites the proof, but raises two questions: how do we find the rank r of the algebra, and (supposing we can answer that) how do we find an element ξ whose degree is r? These are the practical matters we grapple.

Our first contribution is an algorithm for computing the rank of a Euclidean Jordan algebra, thereby answering the first question. The second, we cannot answer per se. Instead we show how Theorem 2 can be reformulated to use an arbitrary basis, alleviating the immediate need to locate a regular element. Both results follow from a characterization in Theorem 3 of the circumstances under which, with respect to an arbitrary basis, the coefficient polynomials $a_0, a_1, \ldots, a_{r-1}$ can be found.

We have essentially already computed a "characteristic polynomial of" function for the Jordan spin algebra in Example 2 with respect to the standard basis. And, working in the standard basis, the need for a regular element never arose. That encouraging example will guide us in the general case.

2 Polynomials and fractions

We adopt the standard interpretation of multivariate polynomials with real coefficients. If $\mathbb{R}[X_1]$ denotes a univariate polynomial ring, then $\mathbb{R}[X_1, X_2]$ is defined recursively as $(\mathbb{R}[X_1])[X_2]$, and so on until we reach $\mathbb{R}[X_1, X_2, \ldots, X_n]$. We will frequently define $R \coloneqq \mathbb{R}[X_1, X_2, \ldots, X_n]$ to keep the notation from getting out of hand.

Each $p \in R$ corresponds to a function $\widehat{p} : \mathbb{R}^n \to \mathbb{R}$ defined in the obvious way so that $\widehat{p}(x_1, x_2, \ldots, x_n)$ is the same expression as p, but with each X_i replaced by x_i. When x is an element of an n-dimensional Euclidean Jordan algebra, we will write $\mathbf{b}(x) \in \mathbb{R}^n$ for the vector of its coordinates with respect to the basis $\mathbf{b}$. Thus if $p \in R$, we can evaluate $\widehat{p}(\mathbf{b}(x)) \in \mathbb{R}$. Similarly, if L is a linear operator, we will write $\mathbf{b}(L)$ for the matrix of L with respect to $\mathbf{b}$ so that $\mathbf{b}(L(x)) = \mathbf{b}(L)\,\mathbf{b}(x)$ for all x.

The multivariate polynomial ring R is an infinite integral domain [1], and the mapping $p \mapsto \widehat{p}$ is a ring isomorphism between R and its image, a ring of functions. Polynomial addition and multiplication are defined precisely so that this map a homomorphism; a recursive root-counting procedure then shows that it is injective [2]. In particular, p is the zero polynomial if and only if $\widehat{p}(x) = 0$ for all $x \in \mathbb{R}^n$. It is a standard exercise in algebraic geometry to show that the complement of every Zariski-closed proper subset of $\mathbb{R}^n$ is both open and dense in the Euclidean topology on $\mathbb{R}^n$. As a result, the set $\{x \in \mathbb{R}^n \mid \widehat{p}(x) \neq 0\}$ is open and dense in $\mathbb{R}^n$ for all nonzero $p \in R$. It is from this raw material that the following result is manufactured.

Proposition 1 (Faraut and Korányi [3], Proposition II.2.1)**.** *The subset of regular elements in any Euclidean Jordan algebra is open and dense in the topology induced by the inner product.*

Since R is an integral domain, its field of fractions $\operatorname{Frac}(R)$ is defined. The set $\operatorname{Frac}(R)$ consists of the equivalence classes $a/b \coloneqq [(a, b)]$ with $b \neq 0_R$ under the relation $a/b \sim c/d \iff ad = bc$ and the latter equality interpreted in R. With the usual kindergarten addition and multiplication defined, the set $\operatorname{Frac}(R)$ forms a field. Its additive identity (zero) element is $0_R/1_R$, and its multiplicative unit element is $1_R/1_R$. The embedding $\iota : R \hookrightarrow \operatorname{Frac}(R)$ with $\iota(a) \coloneqq a/1_R$ is an injective ring homomorphism [1]. In particular, the following are equivalent:

- $a/1_R = \iota(a)$ is zero in $\operatorname{Frac}(R)$

- a is the zero polynomial in R
- $\widehat{a}(x) = 0$ for all $x \in \mathbb{R}^n$.

We will also work with matrices whose entries are polynomials. Any such matrix $A \in R^{m \times n}$ whose entries are $A_{ij} \in R$ corresponds to a function $\widehat{A} : \mathbb{R}^n \to \mathbb{R}^{m \times n}$ such that the entries of $\widehat{A}(x)$ are $\widehat{A_{ij}}(x)$. The entrywise formulae for matrix addition and multiplication show that, when the dimensions of the matrices are compatible, the map $A \mapsto \widehat{A}$ is both additive and multiplicative. We apply this, for example, to conclude that $\widehat{\det(A)}(x) = \det\left(\widehat{A}(x)\right)$ for all $x \in \mathbb{R}^n$. We extend the fraction-field embedding ι to $R^{m \times n}$ by applying it entrywise.

3 Building blocks

We digress to restate explicitly (and in the present notation) two results that underlie our main theorem. They are used implicitly by Faraut and Korányi, but the reader may benefit from seeing them written down.

Recalling that the Jordan product in Definition 1 is bilinear, we may define a linear operator $L_x : V \to V$ on a Euclidean Jordan algebra $(V, \circ, \langle\cdot,\cdot\rangle)$ by $L_x(y) \coloneqq x \circ y$. It follows that $x \mapsto L_x$ is itself linear.

Lemma 1. *If $(V, \circ, \langle\cdot,\cdot\rangle)$ is a Euclidean Jordan algebra of dimension n with basis $\mathbf{b}$ and if $R = \mathbb{R}[X_1, X_2, \ldots, X_n]$, then there exists a matrix $M_{\mathbf{b}} \in R^{n \times n}$ such that*

$$\forall x \in V : \widehat{M_{\mathbf{b}}}(\mathbf{b}(x)) = \mathbf{b}(L_x).$$

Proof. Recall that the map $x \mapsto L_x$ is linear, and let $x = x_1 b_1 + x_2 b_2 + \cdots + x_n b_n$ be the representation of an arbitrary $x \in V$ with respect to the basis $\mathbf{b} = \{b_1, b_2, \ldots, b_n\}$. Then by linearity,

$$L_x = x_1 L_{b_1} + x_2 L_{b_2} + \cdots + x_n L_{b_n},$$

and thus

$$\mathbf{b}(L_x)_{ij} = x_1 \mathbf{b}(L_{b_1})_{ij} + x_2 \mathbf{b}(L_{b_2})_{ij} + \cdots + x_n \mathbf{b}(L_{b_n})_{ij}$$

for all $i, j \in \{1, 2, \ldots, n\}$. Each $\mathbf{b}(L_{b_k})_{ij}$ above is a constant, because the left-multiplication-by-b_k matrix is fixed. It follows that

$$M_{ij} \coloneqq \mathbf{b}(L_{b_1})_{ij} X_1 + \mathbf{b}(L_{b_2})_{ij} X_2 + \cdots + \mathbf{b}(L_{b_n})_{ij} X_n \in R$$

satisfies

$$\forall x \in V : \widehat{M_{ij}}(\mathbf{b}(x)) = \mathbf{b}(L_x)_{ij},$$

because evaluating it at $X_1 = x_1$, $X_2 = x_2$, et cetera, gives us the (i,j)th entry of $\mathbf{b}(L_x)$. All that remains is to define $M_{\mathbf{b}} \coloneqq [M_{ij}]$. □

Notice that in Lemma 1, one set of polynomials works for all elements x of the algebra. As a result, we need only compute the polynomials once, delaying their evaluation until we are given a specific x for which we want to know $\mathbf{b}(L_x)$.

Since $x^k = L_x^k(1_V)$, repeated matrix multiplications can be used to find polynomial column-matrices that produce the $\mathbf{b}$-coordinates of any power x^k. This is the content of the subsequent proposition, whose proof consists more or less of what has already been said.

Proposition 2. *In the context of Lemma 1, define p_0 to be the embedding of $\mathbf{b}(1_V)$ into $R^{n\times 1}$. Then for all $k \in \mathbb{N}$, the polynomial column-matrices $p_k \coloneqq (M_{\mathbf{b}})^k p_0 \in R^{n\times 1}$ satisfy $\widehat{p_k}(\mathbf{b}(x)) = \mathbf{b}(x^k)$.*

Example 3. If $x = (x_1, x_2, x_3)^T$ in the Jordan spin algebra on $\mathbb{R}^3$, then by checking its action on the standard basis $\mathbf{b}$, we see that

$$\mathbf{b}(L_x) = \begin{bmatrix} x_1 & x_2 & x_3 \\ x_2 & x_1 & 0 \\ x_3 & 0 & x_1 \end{bmatrix} \in \mathbb{R}^{3\times 3}.$$

If we let $R \coloneqq \mathbb{R}[X_1, X_2, X_3]$, then as a result, the corresponding polynomial matrix in Lemma 1 is

$$M_{\mathbf{b}} = \begin{bmatrix} X_1 & X_2 & X_3 \\ X_2 & X_1 & 0 \\ X_3 & 0 & X_1 \end{bmatrix} \in R^{3\times 3}.$$

Recalling that the unit element of this algebra is $(1, 0, 0)^T$, we can compute the column matrices in Proposition 2 by applying successive powers of $M_{\mathbf{b}}$,

$$p_0 = \begin{bmatrix} 1_R \\ 0 \\ 0 \end{bmatrix}, \quad p_1 = \begin{bmatrix} X_1 \\ X_2 \\ X_3 \end{bmatrix}, \quad p_2 = \begin{bmatrix} X_1^2 + X_2^2 + X_3^2 \\ 2X_1X_2 \\ 2X_1X_3 \end{bmatrix} \in R^{3\times 1},$$

and so on.

4 Main results

Our two main results involve the solution of polynomial systems over a polynomial fraction field. To show that these solutions have the properties we desire, we will ultimately evaluate the corresponding functions on elements of $\mathbb{R}^n$. This creates a problem, in that we have not said how we intend to turn an element $c/d \in \operatorname{Frac}(\mathbb{R}[X_1, X_2, \ldots, X_n])$ into a function. This turns out to be a subtle issue that is best simply avoided. Instead, we will make heavy use of the fraction field embedding $\iota : R \hookrightarrow \operatorname{Frac}(R)$, using its injectivity to convert back to the ring of polynomials prior to evaluation. This is aesthetically awkward, but our hope is that the didactic dividends offset the price paid in parsimony.

Theorem 3. *Let* $(V, \circ, \langle\cdot,\cdot\rangle)$ *be a Euclidean Jordan algebra of dimension* $n \geq 1$ *and rank* r *with basis* $\mathbf{b}$*. Define* $R \coloneqq \mathbb{R}[X_1, X_2, \ldots, X_n]$ *with its embedding* ι *into* $\mathbb{F} \coloneqq \mathrm{Frac}(R)$*, and let* p_0 *through* $p_n \in R^n$ *be as in Proposition 2. Then if* $s \in \mathbb{N}$ *with* $s \leq r$*, the system*

$$\begin{bmatrix}\iota(p_0) & \iota(p_1) & \cdots & \iota(p_{s-1})\end{bmatrix} a = \iota(p_s) \tag{1}$$

has a solution $a \in \mathbb{F}^{n\times 1}$ *if and only if* $s = r$.

Proof. First we show that $\{\iota(p_0), \iota(p_1), \ldots, \iota(p_{r-1})\}$ is linearly-independent over $\mathbb{F}$. Suppose that

$$\sum_{k=0}^{r-1} \left(\frac{c_k}{d_k}\right) \iota(p_k) = 0_{\mathbb{F}^{n\times 1}}. \tag{2}$$

If all of the c_k are zero, then we are done, so suppose that $\rho \in \mathbb{N}$ is the largest index such that $c_\rho \neq 0_R$. If $\rho = 0$, we contradict ourselves with

$$\left(\frac{c_0}{d_0}\right)\iota(p_0) = 0_{\mathbb{F}^{n\times 1}} \iff c_0 p_0 = 0_{R^{n\times 1}} \iff c_0 = 0_R,$$

since $\widehat{p_0}(z) = \mathbf{b}(1_V) \neq 0_{\mathbb{R}^n}$ for all $z \in \mathbb{R}^n$. Thus we may suppose that $\rho \geq 1$ to avoid empty sums and products in what follows. Cancelling the denominators in Equation (2) and rearranging, we arrive at

$$\left(\prod_{j=0}^{\rho-1} \frac{d_j}{1_R}\right)\left(\frac{c_\rho}{1_R}\right)\iota(p_\rho) + \sum_{k=0}^{\rho-1}\left(\prod_{\substack{j=0\\ j\neq k}}^{\rho} \frac{d_j}{1_R}\right)\left(\frac{c_k}{1_R}\right)\iota(p_k) = 0_{\mathbb{F}^{n\times 1}}.$$

Or, in terms of the fraction-field embedding,

$$\left(\prod_{j=0}^{\rho-1} \iota(d_j)\right)\iota(c_\rho)\,\iota(p_\rho) + \sum_{k=0}^{\rho-1}\left(\prod_{\substack{j=0\\ j\neq k}}^{\rho} \iota(d_j)\right)\iota(c_k)\,\iota(p_k) = \iota(0_{R^{n\times 1}}),$$

which by injectivity reduces to

$$\left(\prod_{j=0}^{\rho-1} d_j\right) c_\rho p_\rho + \sum_{k=0}^{\rho-1}\left(\prod_{\substack{j=0\\ j\neq k}}^{\rho} d_j\right) c_k p_k = 0_{R^{n\times 1}}. \tag{3}$$

We have assumed that $c_\rho \neq 0_R$, and the d_j are all non-zero because they started out as denominators in $\mathbb{F}$. The set of regular elements in V is open and dense—as are the sets where the $\widehat{d_j}$ and $\widehat{c_\rho}$ are nonzero—so we can find a regular element $x \in V$ such that

$$\alpha \coloneqq \widehat{\left(\prod_{j=0}^{\rho-1} d_j\right)}(\mathbf{b}(x))\,\widehat{c_\rho}(\mathbf{b}(x)) \neq 0_{\mathbb{R}}.$$

If we evaluate Equation (3) at $\mathbf{b}(x)$ and divide by this $\alpha \neq 0_{\mathbb{R}}$, then we arrive at an expression of the form

$$\mathbf{b}(x^{\rho}) + \alpha^{-1} \sum_{k=0}^{\rho-1} \beta_k \mathbf{b}\left(x^k\right) = 0_{\mathbb{R}^{n\times 1}}$$

for some collection of $\beta_k \in \mathbb{R}$. Inverting the basis-representation map now gives

$$x^{\rho} + \sum_{k=0}^{\rho-1} \left(\alpha^{-1}\beta_k\right) x^k = 0_V.$$

But this is the result of evaluating a monic univariate polynomial of degree $\rho < r$ at x. Since x is regular, it has degree r, and Definition 2 says that the result cannot be zero for $\rho < r$. We conclude that indeed all of the c_k are zero, and our claim that $\{\iota(p_0), \iota(p_1), \ldots, \iota(p_{r-1})\}$ is linearly-independent follows. This shows that Equation (1) has no solution for $s < r$.

Now if, on the other hand, we have $s = r$, then Equation (1) has a solution. To see this, extend the linearly-independent set $\{\iota(p_0), \iota(p_1), \ldots, \iota(p_{r-1})\}$ to a basis for $\mathbb{F}^n$ by appending elements $\tilde{q}_1, \tilde{q}_2, \ldots, \tilde{q}_{n-r}$. Without loss of generality we assume that each entry of each $\tilde{q}_k$ has denominator one. This can be accomplished without destroying the linear-independence of the set by scaling each $\tilde{q}_k$ by the least common multiple of their denominators, and thus we may presume that these new basis elements satisfy $\tilde{q}_k = \iota(q_k)$ for some $q_k \in R^n$, and that therefore there exists a nonsingular matrix $Q \in R^{n\times n}$ with

$$\iota(Q) \coloneqq \begin{bmatrix} \iota(p_0) & \iota(p_1) & \cdots & \iota(p_{r-1}) & \iota(q_1) & \iota(q_2) & \cdots & \iota(q_{n-r}) \end{bmatrix} \in \mathbb{F}^{n\times n}.$$

Every entry of $\iota(Q)$ has denominator 1_R, and the determinant of $\iota(Q)$ is nonzero because its columns are linearly-independent. As a result, we can apply Cramer's rule to find the unique solution $\tilde{a} = (a_0, a_1, \ldots, a_{n-1})^T$ to the system $\iota(Q)\,\tilde{a} = \iota(p_r)$. If we write $A_{i\leftrightarrow v}$ to denote a matrix A having its ith column replaced by the vector v, then Cramer's rule says that

$$a_i \coloneqq \frac{\det\left(\iota(Q)_{i\leftrightarrow\iota(p_r)}\right)}{\det(\iota(Q))} \in \mathbb{F}.$$

It remains to be seen that $a_i = 0_{\mathbb{F}}$ when $i \geq r$, so that no $\iota(q_k)$ are present in the solution and that therefore $a \coloneqq (a_0, a_1, \ldots, a_{r-1})^T$ solves Equation (1). However, this follows relatively easily from the properties of the determinant. First notice that by definition we have $a_i = 0_{\mathbb{F}}$ if and only if $\det\left(\iota(Q)_{i\leftrightarrow\iota(p_r)}\right) = 0_{\mathbb{F}}$. This determinant is a sum/product of elements of $\mathbb{F}$, so we can apply ι^{-1} to conclude that $a_i = 0_{\mathbb{F}}$ if and only if $\det(Q_{i\leftrightarrow p_r}) = 0_R$. But $\det(Q_{i\leftrightarrow p_r})$ must be zero for $i \geq r$, since the corresponding function from $\mathbb{R}^n \to \mathbb{R}$ is zero on the $\mathbf{b}$-coordinates of any regular element x, the power x^r being a real linear combination of the lower powers in that case. More explicitly,

$$\widehat{\det(Q_{i\leftrightarrow p_r})}(\mathbf{b}(x)) = \det\left(\widehat{Q_{i\leftrightarrow p_r}}(\mathbf{b}(x))\right)$$

and the latter is zero on the dense subset of regular $x \in V$ by the linear-dependence of $\left\{x^0, x^1, \ldots, x^r\right\}$ in that case. By continuity we conclude that $\widehat{\det\left(Q_{i \leftrightarrow p_r}\right)}$ and hence $\det\left(Q_{i \leftrightarrow p_r}\right)$ are zero when $i \geq r$. ☐

This result already provides a (rather wasteful) means to compute the rank of a Euclidean Jordan algebra. Notice that, since rank is bounded above by dimension, we may ignore zero-dimensional algebras entirely. Then to compute the rank of a Euclidean Jordan algebra of dimension $n \geq 1$, we can try to solve Equation (1) for $s = 1, 2, \ldots, n$, stopping when we succeed. The last iteration with $s = n$ can be skipped if we are interested only in the rank and not the solution of the system. If the system isn't solvable for $s < n$, then $s = n$ is the only other possibility for the rank of the algebra. Towards improving this procedure, we notice the following.

Corollary 1. *Let $(V, \circ, \langle \cdot, \cdot \rangle)$ be a Euclidean Jordan algebra of dimension $n \geq 1$ and rank r with basis $\mathbf{b}$. If we define $R \coloneqq \mathbb{R}\left[X_1, X_2, \ldots, X_n\right]$ with its embedding ι into $\mathbb{F} \coloneqq \operatorname{Frac}(R)$ and let p_0 through $p_n \in R^n$ be as in Proposition 2, then the matrix $\begin{bmatrix} \iota(p_0) & \iota(p_1) & \cdots & \iota(p_{n-1}) \end{bmatrix}$ has rank r over $\mathbb{F}$.*

Proof. While proving Theorem 3 we saw that the set $\left\{\iota(p_0), \iota(p_1), \ldots, \iota(p_{r-1})\right\}$ is linearly-independent over $\mathbb{F}$. It follows that the rank of the matrix is at least r and that, moreover, if $r = n$, we are done. This leaves us to prove only that the rank of the matrix cannot exceed r when $r < n$.

With that in mind, take any $k \in \{r, r+1, \ldots, n-1\}$, and let $M_{\mathbf{b}} \in R^{n \times n}$ be the polynomial matrix from Lemma 1. Then, because this is how we constructed p_i for $i \geq 1$,

$$p_k \coloneqq \left(M_{\mathbf{b}}\right)^{k-r} p_r \implies \iota(p_k) = \iota\left(M_{\mathbf{b}}\right)^{k-r} \iota(p_r),$$

and we can use Theorem 3 to replace $\iota(p_r)$ with a linear combination of $\iota(p_0)$ through $\iota(p_{r-1})$,

$$\iota(p_k) = \iota\left(M_{\mathbf{b}}\right)^{k-r} \left(\sum_{i=0}^{r-1} a_i \iota(p_i)\right) = \sum_{i=0}^{r-1} a_i \iota\left(M_{\mathbf{b}}\right)^{k-r} \iota(p_i) = \sum_{i=0}^{r-1} a_i \iota\left(p_{(k-r+i)}\right).$$

Letting $j \coloneqq k - r + i$ in this sum, we see that

$$\iota(p_k) = \sum_{j=k-r}^{k-1} a_{(j-k+r)} \iota(p_j).$$

Thus we have expressed $\iota(p_k)$ in terms of $\iota(p_{k-1})$ through $\iota(p_{k-r})$. If necessary, this process can be repeated to write $\iota(p_{k-1})$ in terms of $\iota(p_{k-2})$ and so forth, until a linear combination involving only $\iota(p_{r-1})$ through $\iota(p_0)$ is reached. This shows that $\iota(p_k) \in \operatorname{span}\left(\left\{\iota(p_0), \iota(p_1), \ldots, \iota(p_{r-1})\right\}\right)$ when $k \geq r$. ☐

With this result at our disposal, we can find the rank of any nontrivial Euclidean Jordan algebra by computing the rank of the matrix in Corollary 1.

Compared to our first idea, this approach is simpler and (if implemented properly) avoids redundant computations. On the other hand, it doesn't necessarily find the coefficients a_i. So unless we are only interested in the rank of the algebra, it also leaves something to be desired.

To emphasize this shortcoming, we present our modified version of Theorem 2 that uses the solution to Equation (1). For contrast, the Faraut and Korányi result is in terms of a basis consisting of powers of a regular element that we know must exist by definition. The regular-element basis is theoretically convenient but practically problematic unless you know how to find a regular element. We allow the customer to bring his own basis. This is above-all advantageous when the standard basis with integer coordinates will suffice and fast rational arithmetic may be used. We omit the proof, which is identical to the one given by Faraut and Korányi after using Theorem 3 to find the coefficients a_i with respect to your favorite basis.

Theorem 4 (BYOB characteristic polynomial of)**.** *If* $(V, \circ, \langle\cdot,\cdot\rangle)$ *is a Euclidean Jordan algebra of rank* r *and dimension* n *with basis* $\mathbf{b}$, *then there exist polynomials* a_0 *through* a_{r-1} *in* $\mathbb{R}[X_1, X_2, \ldots, X_n]$ *such that the characteristic polynomial of any* $x \in V$ *is*

$$\Lambda^r + \sum_{i=0}^{r-1} \widehat{a_i}(\mathbf{b}(x))\,\Lambda^i \in \mathbb{R}[\Lambda].$$

This is nothing more than the Faraut and Korányi result after a change of basis, but its utility stems from the constructive proof it now admits. We point out, just in case it is not clear, that the polynomials a_i in Theorem 4 depend on the basis used. The polynomials p_i in Equation (1) are those obtained in Lemma 1 and Proposition 2 where they were constructed using basis information. As a result, the solution to Equation (1) depends on the basis too.

Our final algorithm we present in more detail. We begin by using Lemma 1 to construct the polynomial matrix $M_{\mathbf{b}}$. We then proceed as in Proposition 2 to compute p_1 through p_n from p_0, which itself can be found using elementary linear algebra. Finally we construct, augment, and row-reduce the matrix from Corollary 1. It is worth keeping in mind that this last step involves nothing more complicated than solving a consistent system over a field.

Algorithm 1 Compute the rank of a Euclidean Jordan algebra V

Input: A basis $\mathbf{b}$ for V and a multiplication table for its Jordan product
Output: The $r := \text{rank}(V)$ polynomials a_0 through a_{r-1} in Theorem 4

```
n ← dim(V) // the cardinality of b
if n = 0 then
    return () // trivial case
end if

Compute the matrices b(L_{b_i}) for all i ∈ {1, 2, ..., n}
for all i, j ∈ {1, 2, ..., n} do
    M_ij ← b(L_{b_1})_ij X_1 + b(L_{b_2})_ij X_2 + ··· + b(L_{b_n})_ij X_n
end for
M_b ← [M_ij] // Lemma 1

Compute b(1_V) from the matrices b(L_{b_i})
p_0 ← the embedding of b(1_V) into R^{n×1}
for all k ∈ {1, 2, ..., n} do
    p_k ← (M_b)^k p_0 // Proposition 2
end for
P ← [ι(p_0)  ι(p_1)  ···  ι(p_{n−1})] // Corollary 1

Augment P with an identity matrix and row-reduce
G ← the augmented portion after reduction
r ← the number of nonzero rows in rref(P)

return the first r entries of −Gp_r
```

Thus we obtain the coefficients a_0 through a_{r-1}, and the rank of the algebra is their number, r. The added efficiency comes at the cost of a bit more complexity and the need to keep track of some row operations.

Example 4. Continuing our running example, let P be the matrix in Corollary 1 that arises from the Jordan spin algebra on $\mathbb{R}^3$. We computed the columns of P in Example 3. Augment P with a three-by-three identity matrix, and call the result P':

$$P' := \left[\begin{array}{ccc|ccc} 1_{\mathbb{F}} & X_1 & X_1^2 + X_2^2 + X_3^2 & 1_{\mathbb{F}} & 0_{\mathbb{F}} & 0_{\mathbb{F}} \\ 0_{\mathbb{F}} & X_2 & 2X_1X_2 & 0_{\mathbb{F}} & 1_{\mathbb{F}} & 0_{\mathbb{F}} \\ 0_{\mathbb{F}} & X_3 & 2X_1X_3 & 0_{\mathbb{F}} & 0_{\mathbb{F}} & 1_{\mathbb{F}} \end{array}\right].$$

After row-reduction, we arrive at

$$\text{rref}(P') = \left[\begin{array}{ccc|ccc} 1_{\mathbb{F}} & 0_{\mathbb{F}} & -X_1^2 + X_2^2 + X_3^2 & 1_{\mathbb{F}} & 0_{\mathbb{F}} & -X_1/X_3 \\ 0_{\mathbb{F}} & 1_{\mathbb{F}} & 2X_1 & 0_{\mathbb{F}} & 0_{\mathbb{F}} & 1/X_3 \\ 0_{\mathbb{F}} & 0_{\mathbb{F}} & 0_{\mathbb{F}} & 0_{\mathbb{F}} & 1_{\mathbb{F}} & -X_2/X_3 \end{array}\right].$$

There are two non-zero rows in the non-augmented portion that corresponds to $\text{rref}(P)$, so $r = 2$. From the theory we know that p_2, the second column of P, is a linear combination of p_0 and p_1. Moreover if G is the augmented portion of the matrix above, then $Pa = p_2$ if and only if $\operatorname{rref}(P)\, a = GPa = Gp_2$, where

$$Gp_2 = \begin{bmatrix} 1_{\mathbb{F}} & 0_{\mathbb{F}} & -X_1/X_3 \\ 0_{\mathbb{F}} & 0_{\mathbb{F}} & 1/X_3 \\ 0_{\mathbb{F}} & 1_{\mathbb{F}} & -X_2/X_3 \end{bmatrix} \begin{bmatrix} X_1^2 + X_2^2 + X_3^2 \\ 2X_1X_2 \\ 2X_1X_3 \end{bmatrix} = \begin{bmatrix} -X_1^2 + X_2^2 + X_3^2 \\ 2X_1 \\ 0_{\mathbb{F}} \end{bmatrix}.$$

Now the system that needs to be solved is $\operatorname{rref}(P)\, a = Gp_2$,

$$\begin{bmatrix} 1_{\mathbb{F}} & 0_{\mathbb{F}} & -X_1^2 + X_2^2 + X_3^2 \\ 0_{\mathbb{F}} & 1_{\mathbb{F}} & 2X_1 \\ 0_{\mathbb{F}} & 0_{\mathbb{F}} & 0_{\mathbb{F}} \end{bmatrix} a = \begin{bmatrix} -X_1^2 + X_2^2 + X_3^2 \\ 2X_1 \\ 0_{\mathbb{F}} \end{bmatrix}.$$

The theory tells us that the first $r = 2$ columns of P were linearly-independent, and that therefore we seek only a_0 and a_1. After pruning the irrelevant rows and columns, we are left with the trivial system,

$$\begin{bmatrix} 1_{\mathbb{F}} & 0_{\mathbb{F}} \\ 0_{\mathbb{F}} & 1_{\mathbb{F}} \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \end{bmatrix} = \begin{bmatrix} -X_1^2 + X_2^2 + X_3^2 \\ 2X_1 \end{bmatrix} \iff \begin{bmatrix} a_0 \\ a_1 \end{bmatrix} = \begin{bmatrix} -X_1^2 + X_2^2 + X_3^2 \\ 2X_1 \end{bmatrix}.$$

Negating the solution to make it fit our definition of the minimal polynomial, we finally arrive at $a_0 = X_1^2 - X_2^2 - X_3^2$ and $a_1 = -2X_1$ with respect to the usual basis in $\mathbb{R}^3$. This agrees (how could it not?) with the solution found in Example 2.

Call our algebra V, and let $e = (1, 1, 1)^T \in V$. With respect to the standard basis $\mathbf{b}$, Theorem 4 says that the characteristic polynomial of e is

$$\Lambda^2 + \widehat{a_1}\left(\mathbf{b}(e)\right)\Lambda^1 + \widehat{a_0}\left(\mathbf{b}(e)\right)\Lambda^0 = \Lambda^2 - 2\Lambda - \Lambda^0.$$

To verify the Cayley-Hamilton theorem for Euclidean Jordan algebras, for example, we may simply replace Λ by e:

$$e^2 - 2e - 1_V = \begin{bmatrix} 3 \\ 2 \\ 2 \end{bmatrix} - 2 \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = 0_V.$$

Small examples build our confidence, but to truly test Algorithm 1, we would prefer a Euclidean Jordan algebra whose rank is not already known. The following is a specific ($m = 2$) member of a family of algebras described in Exercise III.1.b of Faraut and Korányi [3]. In choosing such an example, we face competing interests. If its dimension is too small, the problem is disingenuous: a decomposition into simple algebras is easy to guess. On the other hand, if its dimension is too large, the polynomial computations become unmanageable and must be performed on a computer. We err on the side of authenticity in this case, and must omit some of the computations.

Example 5. Define the block matrix

$$J \coloneqq \begin{bmatrix} 0 & I_2 \\ -I_2 & 0 \end{bmatrix} \in \mathbb{C}^{4\times 4}$$

and consider the *real* vector space

$$V := \left\{ X \in \mathbb{C}^{4\times 4} \;\middle|\; X^T = -X \text{ and } XJ = J\overline{X} \right\},$$

where $\overline{X}$ denotes the entrywise complex conjugate. Under the Jordan- and inner-products $X \circ Y := (XJY + YJX)/2$ and $\langle X, Y \rangle := \operatorname{trace}(X^* Y)$, this vector space forms a Euclidean Jordan algebra. A bit of work shows that $\mathbf{b} = \{b_1, b_2, b_3, b_4, b_5, b_6\}$ forms a basis for V, where

$$b_1 = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},\ b_2 = \begin{bmatrix} 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \end{bmatrix},\ b_3 = \begin{bmatrix} 0 & -i & 0 & 0 \\ i & 0 & 0 & 0 \\ 0 & 0 & 0 & i \\ 0 & 0 & -i & 0 \end{bmatrix},$$

$$b_4 = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{bmatrix},\ b_5 = \begin{bmatrix} 0 & 0 & 0 & -i \\ 0 & 0 & i & 0 \\ 0 & -i & 0 & 0 \\ i & 0 & 0 & 0 \end{bmatrix},\ b_6 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{bmatrix}.$$

Using Lemma 1, we compute

$$M_{\mathbf{b}} = -\frac{1}{2} \begin{bmatrix} 2X_1 & 2X_2 & 2X_3 & 2X_4 & 2X_5 & 0_R \\ X_2 & X_1 + X_6 & 0_R & 0_R & 0_R & X_2 \\ X_3 & 0_R & X_1 + X_6 & 0_R & 0_R & X_3 \\ X_4 & 0_R & 0_R & X_1 + X_6 & 0_R & X_4 \\ X_5 & 0_R & 0_R & 0_R & X_1 + X_6 & X_5 \\ 0_R & 2X_2 & 2X_3 & 2X_4 & 2X_5 & 2X_6 \end{bmatrix}.$$

The unit element in this algebra is $1_V = -J$, and its $\mathbf{b}$-representation is $(-1, 0, 0, 0, 0, -1)^T$. If we embed that vector into $R^{6\times 1}$ and call the result p_0, then the columns p_1 through p_6 can be found using Proposition 2. The results however are not amenable to transcription.

The remainder of Algorithm 1 produces the polynomials $a_0 = X_1 X_6 - X_2^2 - X_3^2 - X_4^2 - X_5^2$ and $a_1 = X_1 + X_6$, showing that the rank of this algebra is two. Moreover by Theorem 4, if $X \in V$ has the basis representation $\mathbf{b}(X) = (x_1, x_2, x_3, x_4, x_5, x_6)^T$, then its characteristic polynomial is

$$\Lambda^2 + (x_1 + x_6)\Lambda + \left(x_1 x_6 - x_2^2 - x_3^2 - x_4^2 - x_5^2\right)\Lambda^0.$$

For a quick check, we note that this gives the correct answers Λ^2 for $X = 0_V$, and $\left(\Lambda - \Lambda^0\right)^2$ for $X = 1_V$.

References

[1] John A. Beachy and William D. Blair. *Abstract Algebra*. Waveland Press, Long Grove, Illinois, fourth ed., 2019. ISBN 9781478638698.

[2] David A. Cox, John Little, and Donal O'Shea. *Ideals, Varieties, and Algorithms*. Undergraduate Texts in Mathematics. Springer, Cham, fourth ed., 2015. ISBN 9783319167206, doi:10.1007/978-3-319-16721-3.

[3] Jacques Faraut and Adam Korányi. *Analysis on Symmetric Cones*. Clarendon Press, Oxford, 1994. ISBN 9780198534778.

[4] Leonid Faybusovich. *Euclidean Jordan algebras and interior-point algorithms*. Positivity, 1(4):331–357, 1997, doi:10.1023/A:1009701824047.

[5] Leonid Faybusovich. *Linear systems in Jordan algebras and primal-dual interior-point algorithms*. Journal of Computational and Applied Mathematics, 86(1):149–175, 1997, doi:10.1016/S0377-0427(97)00153-2.

[6] Leonid Faybusovich. *A Jordan-algebraic approach to potential-reduction algorithms*. Mathematische Zeitschrift, 239(1):117–129, 2002, doi:10.1007/s002090100286.

[7] Osman Güler. *Barrier functions in interior point methods*. Mathematics of Operations Research, 21(4):860–885, 1996, doi:10.1287/moor.21.4.860.

[8] Pascual Jordan, John von Neumann, and Eugene Wigner. *On an algebraic generalization of the quantum mechanical formalism*. Annals of Mathematics, 35(1):29–64, 1934, doi:10.2307/1968117.

[9] Narendra Karmarkar. *A new polynomial-time algorithm for linear programming*. Combinatorica, 4(4):373–395, 1984, doi:10.1007/BF02579150.

[10] Kevin McCrimmon. *A Taste of Jordan Algebras*. Springer-Verlag, New York, 2004. ISBN 0387954473.

[11] Yurii E. Nesterov and Arkadii Nemirovskii. *Interior-Point Polynomial Algorithms in Convex Programming*, vol. 13 of *SIAM Studies in Applied and Numerical Mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, 1994. ISBN 0898713196.