

EXACT SOLUTIONS IN STRUCTURED LOW-RANK APPROXIMATION*

GIORGIO OTTAVIANI[†], PIERRE-JEAN SPAENLEHAUER[‡], AND BERND STURMFELS[§]

Abstract. Structured low-rank approximation is the problem of minimizing a weighted Frobenius distance to a given matrix among all matrices of fixed rank in a linear space of matrices. We study the critical points of this optimization problem using algebraic geometry. A particular focus lies on Hankel matrices, Sylvester matrices, and generic linear spaces.

Key words. low-rank approximation, Gröbner bases, Euclidean distance degree

AMS subject classifications. 14Q15, 65K10, 68W30, 93B11

DOI. 10.1137/13094520X

1. Introduction. *Low-rank approximation* in linear algebra refers to the following optimization problem:

$$(1.1) \quad \text{minimize } \|X - U\|_{\Lambda}^2 = \sum_{i=1}^m \sum_{j=1}^n \lambda_{ij} (x_{ij} - u_{ij})^2 \quad \text{subject to } \text{rank}(X) \leq r.$$

Here, we are given a real *data matrix* $U = (u_{ij})$ of format $m \times n$, and we wish to find a matrix $X = (x_{ij})$ of rank at most r that is closest to U in a weighted Frobenius norm. The entries of the *weight matrix* $\Lambda = (\lambda_{ij})$ are positive reals. If $m \leq n$ and the weight matrix Λ is the all-one matrix $\mathbf{1}$, then the solution to (1.1) is given by the singular value decomposition

$$U = T_1 \cdot \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_m) \cdot T_2.$$

Here T_1, T_2 are orthogonal matrices, and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_m$ are the singular values of U . By the Eckart–Young theorem, the matrix of rank $\leq r$ closest to U equals

$$(1.2) \quad U^* = T_1 \cdot \text{diag}(\sigma_1, \dots, \sigma_r, 0, \dots, 0) \cdot T_2.$$

For weights Λ , the situation is more complicated, as seen in the studies [17, 22, 24]. In particular, there can be many local minima. We discuss a small instance in Example 2.

In *structured low-rank approximation* [5, 18], we are also given a linear subspace $\mathcal{L} \subset \mathbb{R}^{m \times n}$, typically containing the matrix U . We consider the restricted problem:

$$(1.3) \quad \text{minimize } \|X - U\|_{\Lambda}^2 = \sum_{i=1}^m \sum_{j=1}^n \lambda_{ij} (x_{ij} - u_{ij})^2 \quad \text{subject to } X \in \mathcal{L} \text{ and } \text{rank}(X) \leq r.$$

A best-case scenario for $\Lambda = \mathbf{1}$ is this: if U lies in \mathcal{L} , then so does U^* . This happens for some subspaces \mathcal{L} , including symmetric and circulant matrices, but most subspaces \mathcal{L} do not enjoy this property (cf. [5]). Our problem is difficult even for $\Lambda = \mathbf{1}$.

*Received by the editors November 13, 2013; accepted for publication (in revised form) by D. A. Bini September 5, 2014; published electronically December 11, 2014.

<http://www.siam.org/journals/simax/35-4/94520.html>

[†]Università di Firenze, viale Morgagni 67A, 50134 Firenze, Italy (ottavian@math.unifi.it).

[‡]CARAMEL project, INRIA Nancy Grand-Est, Université de Lorraine, CNRS, LORIA, Nancy, France (pierre-jean.spaelnhauer@inria.fr).

[§]University of California, Berkeley, CA 94720-3840 (bernd@berkeley.edu). This author was supported by the NSF (DMS-0968882).

Most practitioners use *local methods* to solve (1.3). These methods return a local minimum. There are many heuristics for ensuring that a local minimum is in fact a global minimum, but there is never a guarantee that this has been accomplished. Another approach is to set up *sum of squares* (SOS) relaxations, which are then solved with semidefinite programming (cf. [2]). These SOS methods furnish certificates of global optimality whenever the relaxation is exact. While this does happen in many instances, there is no a priori guarantee either.

How then can one reliably find all global optima to a polynomial optimization problem such as (1.3)? Aside from interval arithmetic and domain decomposition techniques, the only sure method we are aware of is to list and examine all the critical points. Algorithms that identify the critical points, notably *Gröbner bases* [8] and *numerical algebraic geometry* [1], find all solutions over the complex numbers and sort out the real solutions after the fact. The number of complex critical points is an intrinsic invariant of an optimization problem, and it is a good indicator of the running time needed to solve that problem exactly. The study of such algebraic degrees is an active area of research, and well-developed results are now available for semidefinite programming [21] and maximum likelihood estimation [4].

The present paper applies this philosophy to structured low-rank approximation. A general degree theory for closest points on algebraic varieties was introduced by Draisma et al. in [6]. Following their approach, our primary task is to compute the number of complex critical points of (1.3). Thus, we seek to find the *Euclidean distance degree* (ED degree) of

$$\mathcal{L}_{\leq r} := \{ X \in \mathcal{L} : \text{rank}(X) \leq r \}.$$

This *determinantal variety* is always regarded as a subvariety of the matrix space $\mathbb{R}^{m \times n}$, and we use the Λ -weighted Euclidean distance coming from $\mathbb{R}^{m \times n}$. We write $\text{EDdegree}_{\Lambda}(\mathcal{L}_{\leq r})$ for the Λ -weighted Euclidean distance degree of the variety $\mathcal{L}_{\leq r}$. Thus $\text{EDdegree}_{\Lambda}(\mathcal{L}_{\leq r})$ is the number of complex critical points of the problem (1.3) for generic data matrices U . The importance of keeping track of the weights Λ was highlighted in [6, Example 3.2] for the seemingly harmless situation when \mathcal{L} is the subspace of all symmetric matrices in $\mathbb{R}^{n \times n}$.

Our initial focus lies on the *unit ED degree*, when $\Lambda = \mathbf{1}$ is the all-one matrix, and on the *generic ED degree*, denoted $\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq r})$, when the weight matrix Λ is generic. Choosing generic weights λ_{ij} ensures that the variety $\mathcal{L}_{\leq r}$ meets the isotropic quadric transversally, and it hence allows us to apply formulas from intersection theory such as [6, Theorem 7.7].

This paper is organized as follows. In section 2 we offer a computational study of our optimization problem (1.3) when the subspace \mathcal{L} is generic of codimension c . Two cases are to be distinguished: either \mathcal{L} is a vector space, defined by c homogeneous linear equations in the matrix entries, or \mathcal{L} is an affine space, defined by c inhomogeneous linear equations. We refer to these as the *linear case* and *affine case*, respectively. We present Gröbner basis methods for computing all complex critical points, and we report on their performance. From the complex critical points, one identifies all real critical points and all local minima.

In section 3 we derive some explicit formulas for $\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq r})$ when \mathcal{L} is generic. We cover the four cases that arise by pairing the affine case and the linear case with either unit weights or generic weights. Here we are using techniques from algebraic geometry, including Chern classes and the analysis of singularities. In section 4, we shift gears and focus on special matrices, namely, Hankel matrices and

Sylvester matrices. Those spaces \mathcal{L} arise naturally from symmetric tensor decompositions and approximate gcd computations. These applications require the use of certain specific weight matrices Λ other than $\mathbf{1}$.

We close the introduction with two examples that illustrate the concepts above.

Example 1. Let $m = n = 3$ and $\mathcal{L} \subset \mathbb{R}^{3 \times 3}$ be the five-dimensional space of Hankel matrices:

$$X = \begin{bmatrix} x_0 & x_1 & x_2 \\ x_1 & x_2 & x_3 \\ x_2 & x_3 & x_4 \end{bmatrix}, \quad U = \begin{bmatrix} u_0 & u_1 & u_2 \\ u_1 & u_2 & u_3 \\ u_2 & u_3 & u_4 \end{bmatrix}, \quad \text{and} \quad \Lambda = \begin{bmatrix} \lambda_0 & \lambda_1 & \lambda_2 \\ \lambda_1 & \lambda_2 & \lambda_3 \\ \lambda_2 & \lambda_3 & \lambda_4 \end{bmatrix}.$$

Our goal in (1.3) is to solve the following constrained optimization problem for $r = 1, 2$:

$$\begin{aligned} &\text{minimize } \lambda_0(x_0 - u_0)^2 + 2\lambda_1(x_1 - u_1)^2 \\ &\quad + 3\lambda_2(x_2 - u_2)^2 + 2\lambda_3(x_3 - u_3)^2 + \lambda_4(x_4 - u_4)^2 \\ &\text{subject to } \text{rank}(X) \leq r. \end{aligned}$$

This can be stated as an unconstrained optimization problem. For instance, for rank $r = 1$, we get a one-to-one parametrization of $\mathcal{L}_{\leq 1}$ by setting $x_i = st^i$, and we seek to

$$\text{minimize } \lambda_0(s - u_0)^2 + 2\lambda_1(st - u_1)^2 + 3\lambda_2(st^2 - u_2)^2 + 2\lambda_3(st^3 - u_3)^2 + \lambda_4(st^4 - u_4)^2.$$

The ED degree is the number of critical points with $t \neq 0$. We consider three weights:

$$\mathbf{1} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}, \quad \Omega = \begin{bmatrix} 1 & 1/2 & 1/3 \\ 1/2 & 1/3 & 1/2 \\ 1/3 & 1/2 & 1 \end{bmatrix}, \quad \Theta = \begin{bmatrix} 1 & 2 & 2 \\ 2 & 2 & 2 \\ 2 & 2 & 1 \end{bmatrix}.$$

Here Ω gives the usual Euclidean metric when \mathcal{L} is identified with \mathbb{R}^5 , and Θ arises from identifying \mathcal{L} with symmetric $2 \times 2 \times 2$ -tensors, as in section 4. We compute

$$\begin{aligned} \text{EDdegree}_{\mathbf{1}}(\mathcal{L}_{\leq 1}) &= 6, & \text{EDdegree}_{\Omega}(\mathcal{L}_{\leq 1}) &= 10, & \text{EDdegree}_{\Theta}(\mathcal{L}_{\leq 1}) &= 4, \\ \text{EDdegree}_{\mathbf{1}}(\mathcal{L}_{\leq 2}) &= 9, & \text{EDdegree}_{\Omega}(\mathcal{L}_{\leq 2}) &= 13, & \text{EDdegree}_{\Theta}(\mathcal{L}_{\leq 2}) &= 7. \end{aligned}$$

In both cases, Ω exhibits the generic behavior: $\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq r}) = \text{EDdegree}_{\Omega}(\mathcal{L}_{\leq r})$. See sections 3 and 4 for larger Hankel matrices and formulas for their ED degrees. \square

Example 2. Let $m = n = 3, r = 1$, but now take $\mathcal{L} = \mathbb{R}^{3 \times 3}$, so this is just the weighted rank-one approximation problem for 3×3 -matrices. We know from [6, Example 7.10] that $\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq 1}) = 39$. We take a circulant data matrix and a circulant weight matrix:

$$U = \begin{bmatrix} -59 & 11 & 59 \\ 11 & 59 & -59 \\ 59 & -59 & 11 \end{bmatrix} \quad \text{and} \quad \Lambda = \begin{bmatrix} 9 & 6 & 1 \\ 6 & 1 & 9 \\ 1 & 9 & 6 \end{bmatrix}.$$

This instance has 39 critical points. Of these, 19 are real, and 7 are local minima:

$$\begin{aligned} & \begin{bmatrix} 0.0826 & 2.7921 & -1.5452 \\ 2.7921 & 94.3235 & -52.2007 \\ -1.5452 & -52.2007 & 28.8890 \end{bmatrix}, \begin{bmatrix} -52.2007 & 28.8890 & -1.5452 \\ 2.7921 & -1.5452 & 0.0826 \\ 94.3235 & -52.2007 & 2.7921 \end{bmatrix}, \\ & \begin{bmatrix} -52.2007 & 2.7921 & 94.3235 \\ 28.8890 & -1.5452 & -52.2007 \\ -1.5452 & 0.0826 & 2.7921 \end{bmatrix}, \begin{bmatrix} -29.8794 & 36.2165 & -27.2599 \\ -32.7508 & 39.6968 & -29.8794 \\ 39.6968 & -48.1160 & 36.2165 \end{bmatrix}, \\ & \begin{bmatrix} -48.1160 & 36.2165 & 39.6968 \\ 36.2165 & -27.2599 & -29.8794 \\ 39.6968 & -29.8794 & -32.7508 \end{bmatrix}, \begin{bmatrix} -29.8794 & -32.7508 & 39.6968 \\ 36.2165 & 39.6968 & -48.1160 \\ -27.2599 & -29.8794 & 36.2165 \end{bmatrix}, \\ & \begin{bmatrix} -25.375 & -25.375 & -25.375 \\ -25.375 & -25.375 & -25.375 \\ -25.375 & -25.375 & -25.375 \end{bmatrix}. \end{aligned}$$

The first three are the global minima. The last matrix is the local minimum where the objective function has the largest value: note that each entry equals $-203/8$. The entries of the first six matrices are algebraic numbers of degree 10 over \mathbb{Q} . For instance, the two upper left entries 0.0826 and -48.1160 are among the four real roots of the irreducible polynomial

$$\begin{aligned} & 164466028468224x^{10} + 27858648335954688x^9 + 1602205386689376672x^8 \\ & + 7285836260028875412x^7 - 2198728936046680414272x^6 \\ & - 14854532690380098143152x^5 + 2688673091228371095762316x^4 \\ & + 44612094455115888622678587x^3 - 41350080445712457319337106x^2 \\ & + 27039129499043116889674775x - 1977632463563766878765625. \end{aligned}$$

Thus, the critical ideal in $\mathbb{Q}[x_{11}, x_{12}, \dots, x_{33}]$ is not prime. It is the intersection of six maximal ideals. Their degrees over \mathbb{Q} are 1, 2, 6, 10, 10, 10, for a total of $39 = \text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq 1})$. \square

Rey [22] reports on numerical experiments with the optimization problem (1.1), and he asks whether the number of local minima is bounded above by $\min(m, n)$. Our Example 2 gives a negative answer: the number of local minima can exceed $\min(m, n)$. This result highlights the value of our exact algebraic methods for practitioners of optimization.

2. Gröbner bases. The critical points of the low-rank approximation problem (1.3) can be computed as the solution set of a system of polynomial equations. In this section we derive these equations, and we demonstrate how to solve a range of instances using current Gröbner basis techniques. Here, our emphasis lies on the case when \mathcal{L} is a generic subspace, either linear or affine.

Starting with the linear case, let $\{L_1, L_2, \dots, L_s\}$ be a basis of \mathcal{L}^\perp , the space of linear forms on $\mathbb{R}^{m \times n}$ that vanish on \mathcal{L} . Thus $\text{codim}(\mathcal{L}) = s$, each derivative $\partial L_k / \partial x_{ij}$ is a constant, and $\mathcal{L} = \{X \in \mathbb{R}^{m \times n} : L_1(X) = \dots = L_s(X) = 0\}$. The case when \mathcal{L} is an affine space can be treated with the same notation if we take each L_i to be a linear form plus a constant.

The following implicit formulation of the critical equations is a variation on [6, (2.1)]. We begin with the case $m = n = r + 1$. Let $D \in \mathbb{Z}[x_{11}, \dots, x_{nn}]$ denote the determinant of the $n \times n$ -matrix $X = (x_{ij})$. Given a data matrix $U = (u_{ij}) \in \mathbb{R}^{n \times n}$,

the critical points of $\sum_{i=1}^n \sum_{j=1}^n \lambda_{ij}(x_{ij} - u_{ij})^2$ on the determinantal hypersurface $\mathcal{L}_{\leq n-1} = \{X \in \mathcal{L} : D(X) = 0\}$ verify the following conditions. The matrix on the right has $s + 2$ rows and n^2 columns:

$$\begin{cases} D(X) = 0 \\ L_1(X) = 0 \\ \vdots \\ L_s(X) = 0 \end{cases} \quad \text{Rank} \begin{bmatrix} \partial D/\partial x_{11} & \cdots & \partial D/\partial x_{nn} \\ \partial L_1/\partial x_{11} & \cdots & \partial L_1/\partial x_{nn} \\ \vdots & \ddots & \vdots \\ \partial L_s/\partial x_{11} & \cdots & \partial L_s/\partial x_{nn} \\ \lambda_{11}(x_{11} - u_{11}) & \cdots & \lambda_{nn}(x_{nn} - u_{nn}) \end{bmatrix} \leq s + 1.$$

Any singular point of $\mathcal{L}_{\leq n-1}$ also satisfies these conditions. The rank condition on the Jacobian matrix can be modeled by introducing Lagrange multipliers z_0, z_1, \dots, z_s . These are new variables. We now consider the following polynomial system in $n^2 + s + 1$ variables:

$$(2.1) \quad \begin{cases} D(X) = 0 \\ L_1(X) = 0 \\ \vdots \\ L_s(X) = 0 \end{cases} \quad [z_0 \cdots z_s \ 1] \cdot \begin{bmatrix} \partial D/\partial x_{11} & \cdots & \partial D/\partial x_{nn} \\ \partial L_1/\partial x_{11} & \cdots & \partial L_1/\partial x_{nn} \\ \vdots & \ddots & \vdots \\ \partial L_s/\partial x_{11} & \cdots & \partial L_s/\partial x_{nn} \\ \lambda_{11}(x_{11} - u_{11}) & \cdots & \lambda_{nn}(x_{nn} - u_{nn}) \end{bmatrix} = [0 \cdots 0].$$

Table 1 shows the number of complex solutions to these equations. These numbers are obtained from the formulas in section 3. We verified them using Gröbner bases.

We observe that Table 1 has the following remarkable properties:

- There is a shift between the ED degrees of affine and linear sections for $s \geq (n - 1)^2$. This phenomenon will be explained in Proposition 3.1.
- For Λ general, the third block of columns (linear entries) is constant for $s \leq n(n - 2)$ and the fourth one (affine entries) is constant for $s \leq n(n - 2) + 1$. This is explained in Corollaries 3.2 and 3.5.
- The differences between the first and the third block of columns (both with linear entries) equal those between the second and the fourth one (both with affine entries). This gap is expressed (conjecturally) with formula (3.6).

We prove the correctness of the formulation (2.1) and then discuss our computations.

PROPOSITION 2.1. *For a generic linear (or affine) space \mathcal{L} of codimension s and for a generic data matrix $U = (u_{ij})$ in \mathcal{L} , the solutions (X, z) of the polynomial system (2.1) correspond to the critical points X of the optimization problem (1.3) for square matrices of corank one.*

Proof. We prove this for linear spaces \mathcal{L} . The argument is similar when \mathcal{L} is an affine space. Any solution of the system (2.1) corresponds to a point of \mathcal{L} where the Jacobian matrix of $(D, L_1, \dots, L_s, \|X - U\|_\Lambda^2)$ has a rank defect. There are two types of such points: the critical points of the distance function and singular points on the determinantal variety. Hence it suffices to prove that no point in the singular locus corresponds to a solution of (2.1). The matrix $U = (u_{ij})_{1 \leq i, j \leq n}$ was assumed to be generic, so it has rank n since \mathcal{L} is also generic.

If X is a singular point of the linear section of the variety defined by $D(X) = L_1(X) = \dots = L_s(X) = 0$, then there exists (y_0, y_1, \dots, y_s) with $y_0 \neq 0$ such that

$$[y_0 \ y_1 \ \cdots \ y_s] \cdot \begin{bmatrix} \partial D/\partial x_{11} & \cdots & \partial D/\partial x_{nn} \\ \partial L_1/\partial x_{11} & \cdots & \partial L_1/\partial x_{nn} \\ \vdots & \ddots & \vdots \\ \partial L_s/\partial x_{11} & \cdots & \partial L_s/\partial x_{nn} \end{bmatrix} = [0 \ \cdots \ 0].$$

TABLE 1
 The ED degree for the determinant of an $n \times n$ -matrix with linear or affine entries.

Linear, $\Lambda = \mathbf{1}$					Affine, $\Lambda = \mathbf{1}$			
$n =$	2	3	4	5	2	3	4	5
$s = 0$	2	3	4	5	2	3	4	5
$s = 1$	4	15	28	45	6	15	28	45
$s = 2$	2	31	92	205	4	31	92	205
$s = 3$	0	39	188	605	2	39	188	605
$s = 4$		33	260	1221		39	260	1221
$s = 5$		21	284	1805		33	284	1805
$s = 6$		9	284	2125		21	284	2125
$s = 7$		3	284	2205		9	284	2205
$s = 8$		0	284	2205		3	284	2205
$s = 9$			264	2205			284	2205
$s = 10$			204	2205			264	2205
$s = 11$			120	2205			204	2205
$s = 12$			52	2205			120	2205
$s = 13$			16	2205			52	2205
$s = 14$			4	2205			16	2205
$s = 15$			0	2205			4	2205

Linear, Λ gen.					Affine, Λ gen.			
$n =$	2	3	4	5	2	3	4	5
$s = 0$	6	39	284	2205	6	39	284	2205
$s = 1$	4	39	284	2205	6	39	284	2205
$s = 2$	2	39	284	2205	4	39	284	2205
$s = 3$	0	39	284	2205	2	39	284	2205
$s = 4$		33	284	2205		39	284	2205
$s = 5$		21	284	2205		33	284	2205
$s = 6$		9	284	2205		21	284	2205
$s = 7$		3	284	2205		9	284	2205
$s = 8$		0	284	2205		3	284	2205
$s = 9$			264	2205			284	2205
$s = 10$			204	2205			264	2205
$s = 11$			120	2205			204	2205
$s = 12$			52	2205			120	2205
$s = 13$			16	2205			52	2205
$s = 14$			4	2205			16	2205
$s = 15$			0	2205			4	2205

Let us assume by contradiction that X extends to a solution (X, z) of (2.1). Then

$$\begin{aligned}
 & \left[0 \left(z_1 - \frac{y_1 z_0}{y_0} \right) \cdots \left(z_s - \frac{y_s z_0}{y_0} \right) 1 \right] \cdot \begin{bmatrix} \partial D / \partial x_{11} & \cdots & \partial D / \partial x_{nn} \\ \partial L_1 / \partial x_{11} & \cdots & \partial L_1 / \partial x_{nn} \\ \vdots & \ddots & \vdots \\ \partial L_s / \partial x_{11} & \cdots & \partial L_s / \partial x_{nn} \\ \lambda_{11}(x_{11} - u_{11}) & \cdots & \lambda_{nn}(x_{nn} - u_{nn}) \end{bmatrix} \\
 & = [0 \cdots 0].
 \end{aligned}$$

This means that $X - U$ belongs to \mathcal{L} , and $\Lambda * (X - U)$ belongs to \mathcal{L}^\perp . Here $*$ denotes the Hadamard (coordinatewise) product of two matrices. The scalar product of $X - U$ and $\Lambda * (X - U)$ is zero. Since all coordinates live in \mathbb{R} , these conditions imply $\|X - U\|_\Lambda^2 = 0$, and hence $X = U$. We get a contradiction since U has full rank, whereas $D(X) = 0$. \square

The values of $\text{EDdegree}_\Lambda(\mathcal{L}_{\leq n-1})$ in Table 1 can be verified computationally with the formulation (2.1). We used the implementation of Faugère’s Gröbner basis algorithm F_5 [8] in the `maple` package `FGb`. Computing Gröbner bases for (2.1) was fairly easy for $n \leq 4$ but difficult already for $n = 5$. For each of the cases in Table 1, we computed the ED degree by running `FGb` over the finite field with 65521 elements. However, due to substantial coefficient growth, this did not work over the field \mathbb{Q} of rational numbers. Hence, to actually compute all critical points over \mathbb{C} and hence all local minima over \mathbb{R} , even for $n = 4$, a better formulation was required. In what follows we shall present two such improved formulations.

Duality plays a key role in the computation of the critical points of the Euclidean distance and was investigated in [6, section 5]. In what follows, we compute the critical points of the weighted Euclidean distance of the determinant by using this duality. In the following statement we are using the standing hypothesis that all λ_{ij} are nonzero.

PROPOSITION 2.2. *Let U be a generic $m \times n$ matrix with $m \leq n$, let Λ be a weight matrix, and fix an integer $r \leq \min(m, n)$. Then there is a bijection between the critical points of*

- (1) $Q(X) = \sum_{i,j} \lambda_{ij} (x_{ij} - u_{ij})^2$ on the variety $\mathbb{C}_{\leq m-r}^{m \times n}$ of corank r matrices X , and
- (2) $Q_{\text{dual}}(Y) = \sum_{i,j} (y_{ij} - \lambda_{ij} u_{ij})^2 / \lambda_{ij}$ on the variety $\mathbb{C}_{\leq r}^{m \times n}$ of rank r matrices Y .

For each critical point X of (1), the corresponding critical point Y of (2) equals $Y = \Lambda * U - \Lambda * X$, where $*$ denotes the Hadamard product. In particular, if U has real entries, then the bijection interchanges the real critical points of (1) and of (2).

Proof. The critical points of (1) correspond to matrices X such that the Hadamard product $\Lambda * (U - X)$ is perpendicular to the tangent space at X of the variety $\mathbb{C}_{\leq m-r}^{m \times n}$ of corank r matrices. Recall, e.g., from [6, section 5], that the dual variety to $\mathbb{C}_{\leq m-r}^{m \times n}$ is the variety $\mathbb{C}_{\leq r}^{m \times n}$ of rank r matrices. Hence, the critical points in (1) can be found by solving the linear equation $Y = \Lambda * (U - X)$ on the conormal variety. That *conormal variety* is the set of all pairs (X, Y) such that $X \in \mathbb{C}_{\leq m-r}^{m \times n}$, $Y \in \mathbb{C}_{\leq r}^{m \times n}$, $X^t \cdot Y = 0$, and $X \cdot Y^t = 0$. We can now express X in terms of Y and the parameters by writing $X = \Lambda^{*-1} * (\Lambda * U - Y)$, where Λ^{*-1} denotes the Hadamard (coordinatewise) inverse of the weight matrix Λ . Using biduality, this means that $X = \Lambda^{*-1} * (\Lambda * U - Y)$ is perpendicular to the tangent space at Y of the variety $\mathbb{C}_{\leq r}^{m \times n}$. This is equivalent to the statement that Y is a critical point of (2) on $\mathbb{C}_{\leq r}^{m \times n}$. \square

In both Propositions 2.1 and 2.2, it is assumed that the given matrix U is generic. Here the term *generic* is meant in the usual sense of algebraic geometry: U lies in the complement of an algebraic hypersurface. In particular, that complement is dense in $\mathbb{R}^{m \times n}$, so U will be generic with probability one when drawn from a probability measure supported on $\mathbb{R}^{m \times n}$. However, an exact characterization of genericity is difficult. The polynomial that defines the aforementioned hypersurface is the *ED discriminant*. As can be seen in [6, section 7], this is a very large polynomial of high degree, and we will rarely be able to identify it in an explicit way.

Proposition 2.2 shows that weighted low-rank approximation can be solved by the dual problem. We focus now on the corank 1 case (whose dual problem is rank 1 approximation). For this, we use the parametrization of $n \times n$ matrices of rank 1 by

$$(2.2) \quad (t_1, \dots, t_n, z_1, \dots, z_{n-1}) \mapsto \begin{bmatrix} t_1 & t_1 z_1 & \dots & t_1 z_{n-1} \\ \vdots & \vdots & \vdots & \vdots \\ t_n & t_n z_1 & \dots & t_n z_{n-1} \end{bmatrix}.$$

Remark 1. This parametrization is not surjective: the rank 1 matrices whose first column is zero are missing. This is not an issue when U and \mathcal{L} are generic, since in that case all critical points are in the image of the parametrization. However, for specific U or \mathcal{L} , if some of the critical points are missing, they can be computed by choosing n such parametrizations whose ranges cover all rank 1 matrices. This multiplies the computation time by n . Our a priori computation of the ED degree is also useful to overcome these difficulties. Suppose that the expected number of critical points is known. Then, after some parametrizations have been tried for the given data (U, \mathcal{L}) , the user is guaranteed that all critical points have been found.

The parametrization (2.2) expresses the dual problem (for corank one) as an unconstrained optimization problem in $2n - 1$ variables:

$$(2.3) \text{ Maximize } Q_{\text{dual}} = \sum_{1 \leq i, j \leq n} \frac{1}{\lambda_{ij}} (y_{ij} - \lambda_{ij} u_{ij})^2, \text{ where } y_{i1} = t_i \text{ and } y_{ij} = t_i z_{j-1}.$$

Here, “maximize” is used in an unconventional way: what we seek is the critical point furthest to U . That critical point need not be a local maximum; see, e.g., [6, Figure 4]. We compute the critical points for (2.3) by applying Gröbner bases to the equations

$$\partial Q_{\text{dual}} / \partial t_i = \partial Q_{\text{dual}} / \partial z_j = 0 \quad \text{for } i \in \{1, \dots, n\} \text{ and } j \in \{1, \dots, n - 1\}.$$

The critical points of the primal problem are found by the formula $Y = \Lambda * (U - X)$.

This concludes our discussion of square matrices of rank 1 or corank 1. We next consider the general case of rectangular matrices of format $m \times n$ with general linear or affine entries. We assume $r \leq m \leq n$ and $s \leq mn$. Let M be a complex $m \times n$ -matrix of rank r . Then M is a smooth point in the variety $\mathbb{C}_{\leq r}^{m \times n}$ of matrices of rank $\leq r$. Let $\text{Ker}_L(M)$ and $\text{Ker}_R(M)$ denote the left and right kernels of M , respectively. The normal space of $\mathbb{C}_{\leq r}^{m \times n}$ at M has dimension $(m - r)(n - r)$, and it equals $\text{Ker}_L(M) \otimes \text{Ker}_R(M) \subset \mathbb{C}^{m \times n}$ [11, Chapter 6]. Its orthogonal complement is the tangent space at M , which has dimension $rm + rn - r^2$.

In order to construct a polynomial system whose solutions are the critical points of $X \mapsto \|X - U\|_{\Lambda}^2$ on the smooth locus of $\mathcal{L}_{\leq r}$, we introduce two matrices of unknowns:

$$Y = \begin{bmatrix} 1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 1 \\ y_{1,1} & \dots & y_{1,m-r} \\ \vdots & \ddots & \vdots \\ y_{r,1} & \dots & y_{r,m-r} \end{bmatrix} \quad \text{and} \quad Z = \begin{bmatrix} 1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 1 \\ z_{1,1} & \dots & z_{1,n-r} \\ \vdots & \ddots & \vdots \\ z_{r,1} & \dots & z_{r,n-r} \end{bmatrix}.$$

For $i \in \{1, \dots, m - r\}$, $j \in \{1, \dots, n - r\}$, let $N^{((m-r)(j-1)+i)}$ be the rank 1 matrix which is the product of the i th column of Y and of the j th row of Z^T . We consider

$$(2.4) \quad \begin{cases} Y^T \cdot X = 0 \\ X \cdot Z = 0 \\ L_1(X) = 0 \\ \vdots \\ L_s(X) = 0 \end{cases} \quad [w_1 \ \dots \ w_{(m-r)(n-r)+s} \ 1] \quad \begin{bmatrix} N_{11}^{(1)} & \dots & N_{mn}^{(1)} \\ \vdots & \ddots & \vdots \\ N_{11}^{((m-r)(n-r))} & \dots & N_{mn}^{((m-r)(n-r))} \\ \partial L_1 / \partial x_{11} & \dots & \partial L_1 / \partial x_{mn} \\ \vdots & \ddots & \vdots \\ \partial L_s / \partial x_{11} & \dots & \partial L_s / \partial x_{mn} \\ \lambda_{11}(x_{11} - u_{11}) & \dots & \lambda_{mn}(x_{mn} - u_{mn}) \end{bmatrix} = 0.$$

TABLE 2

Symbolic computation of the weighted rank 3 approximations of a 4×4 matrix.

	Determinant primal (2.1)	Parametric dual (2.3)	Normal space primal (2.4)	Normal space dual (2.4)
Λ generic, GF(65521)	5s	1.3s	6s	8.6s
Λ generic, over \mathbb{Q}	> 1 day	891s	1327s	927s
$\Lambda = \mathbf{1}$, over \mathbb{Q}	0.3s	0.2s	0.4s	0.5s

The rank condition on the matrix in (2.4) comes from the fact that $M \in \mathcal{L}_{\leq r}$ is a critical point if the gradient of the distance function at M belongs to the normal space of $\mathcal{L}_{\leq r}$ at M . The first $(m-r)(n-r) + s$ rows of the matrix span the normal space of $\mathcal{L}_{\leq r}$ at a smooth point. This formulation avoids saturating by the singular locus, which is often too costly.

PROPOSITION 2.3. *For a generic affine space \mathcal{L} of codimension s and a generic matrix U in \mathcal{L} , the polynomial system 2.4 has finitely many complex solutions which correspond to the critical points of the weighted Euclidean distance function on the smooth locus of $\mathcal{L}_{\leq r}$.*

Proof. This is derived from [6, Lemma 2.1]. It is analogous to Proposition 2.1. \square

As in the corank 1 case, for special data (U, \mathcal{L}) some critical points may be missed because our formulation computes only the critical points in a dense open subset of $\mathcal{L}_{\leq r}$. However, the same fix as in Remark 1 works here. We can redo the computations in any of the $\binom{n}{r} \binom{m}{r}$ charts corresponding to the invertibility of pairs of square submatrices of Y and Z .

We next discuss our computational experience with Gröbner bases. In Table 2, we compare the efficiency of the different approaches on a specific problem: computing the weighted rank 3 approximation of a 4×4 matrix. The experimental setting is the following: we consider a 4×4 matrix U with integer entries picked uniformly at random in $\{-100, \dots, 100\}$ and a random weight matrix Λ with positive integer entries chosen at random in $\{1, \dots, 20\}$. By Table 1, the generic ED degree is 284 and the ED degree for $\Lambda = \mathbf{1}$ is 4. We report in Table 2 the timings for computing a lexicographical Gröbner basis with the `maple` package `FGb` [8]. Once a Gröbner basis is known, isolation techniques may be used to obtain the real roots. The `maple` package `fgbrs` provides implementations of such methods.

We examine three scenarios. In the first row, the computation is performed over a finite field. This gives information about the algebraic difficulty of the problem: there is no coefficient growth, and the timings indicate the number of arithmetic operations in Gröbner bases algorithms. However, finding local minima requires computing over \mathbb{Q} . In rows 2 and 3 of Table 2, we compare the case of generic weights with the unweighted case (1.2) that corresponds to the singular value decomposition ($\Lambda = \mathbf{1}$). The dual problem is easiest to solve, in particular with the unconstrained formulation (2.3). Note that, for $s \geq 1$, such an unconstrained formulation is not available, since $\mathcal{L}_{\leq r}$ is generally not a unirational variety.

In Table 3, we report on some Gröbner basis computations with the `maple` package `FGb` for $\Lambda = \mathbf{1}$. Here we used the formulation (2.4). The ED degree, given in bold face, is followed by the time, measured in seconds, for computing the graded reverse lexicographic Gröbner basis. The first timing is obtained by performing the computation over the finite field GF(65521); the second one is obtained by computing over the field of rationals \mathbb{Q} . The symbol “—” means that we did not obtain the Gröbner basis after seven days of computation.

TABLE 3
Symbolic computations for affine sections of determinantal varieties with $\Lambda = \mathbf{1}$.

(m, n, r)	$s = 0$	$s = 1$	$s = 2$	$s = 3$
$(4, 4, 2)$	4/0.42s/1.8s	54/1.93s/744s	230/52.7s/-	582/349.2s/-
$(3, 4, 2)$	3/0.2s/0.3s	15/0.3s/7.4s	43/1.2s/132s	71/1.5s/1120s
$(3, 5, 2)$	3/0.3s/0.5s	15/0.5s/16s	43/2.1s/400s	87/7.1s/6038s

(m, n, r)	$s = 4$	$s = 5$	$s = 6$
$(4, 4, 2)$	998/1474s/-	1250/2739s/-	1250/2961s/-
$(3, 4, 2)$	83/2.2s/2696s	83/2.3s/4846s	83/2.1s/5764s
$(3, 5, 2)$	127/16s/59091s	143/20s/160094s	143/20s/68164s

(m, n, r)	$s = 7$	$s = 8$	$s = 9$
$(4, 4, 2)$	1074/1816s/-	818/821s/-	532/349s/-
$(3, 4, 2)$	73/2.2s/4570s	49/1.0s/1619s	22/0.8s/350s
$(3, 5, 2)$	143/20s/99208s	143/20s/163532s	128/18s/263586s

(m, n, r)	$s = 10$	$s = 11$	$s = 12$
$(4, 4, 2)$	276/92s/-	100/42s/450988s	20/1.4s/1970s
$(3, 4, 2)$	6/0.3s/6.4s		
$(3, 5, 2)$	88/13s/67460s	40/1.9s/4568s	10/0.8s/114s

An important observation in Table 3 is the correlation between the reported running times and the values of EDdegree_1 . The former tell us how many arithmetic operations are needed to find a Gröbner basis. This suggests that the ED degree is an accurate measure for the complexity of solving low-rank approximation problems with symbolic algorithms, and it serves as a key motivation for computing ED degrees using advanced tools from algebraic geometry. This will be carried out in the next section, both for Λ generic and for $\Lambda = \mathbf{1}$. In particular, we shall arrive at theoretical explanations for the ED degrees in Tables 1 and 3.

3. Algebraic geometry. The study of ED degrees for algebraic varieties was started in [6]. This section builds on and further develops the geometric theory in that paper. We focus on the low rank approximation problem (1.3), and we derive general formulas for the ED degrees in Tables 1 and 3.

We recall that an affine variety $X \subset \mathbb{C}^{N+1}$ is an *affine cone* if $x \in X$ implies $tx \in X$ for every $t \in \mathbb{C}$. The variety of $m \times n$ -matrices of rank $\leq r$ is an affine cone. If $X \subset \mathbb{C}^{N+1}$ is an affine cone, then the corresponding projective variety $\mathbb{P}X \subset \mathbb{P}^N$ is well defined. The ED degree of $\mathbb{P}X$ is the ED degree of its affine cone X . The following proposition explains the shift between the third and fourth column of Table 1. More generally, it shows that we can restrict the analysis to linear sections, since the ED degree (for generic weights) in the affine case can be deduced from the linear case.

PROPOSITION 3.1. *Let $X \subset \mathbb{C}^{N+1}$ be an affine cone, let \mathcal{A}^s (resp., \mathcal{L}^s) be a generic affine (resp., linear) subspace of codimension $s \geq 1$ in \mathbb{C}^{N+1} . Then*

$$(3.1) \quad \text{EDdegree}_{\text{gen}}(X \cap \mathcal{A}^s) = \text{EDdegree}_{\text{gen}}(X \cap \mathcal{L}^{s-1}).$$

Proof. Let $\overline{X} \subset \mathbb{P}^{N+1}$ be the projective closure of X . From [6, Theorem 6.11], we have $\text{EDdegree}_{\text{gen}}(X) = \text{EDdegree}_{\text{gen}}(\overline{X})$, since the transversality assumptions in that result are satisfied for general weights. From the equality $\overline{X \cap \mathcal{A}^s} = \overline{X} \cap \mathcal{L}^s$, we conclude $\text{EDdegree}_{\text{gen}}(X \cap \mathcal{A}^s) = \text{EDdegree}_{\text{gen}}(\overline{X} \cap \mathcal{L}^s) = \text{EDdegree}_{\text{gen}}(\mathbb{P}X \cap \mathcal{L}^{s-1}) = \text{EDdegree}_{\text{gen}}(X \cap \mathcal{L}^{s-1})$. Here, the second equality follows from $\mathbb{P}X = \overline{X} \cap \mathcal{L}^1$. \square

Consider a projective variety X embedded in \mathbb{P}^N with a generic system of coordinates. It was shown in [6, Theorem 5.4] that $\text{EDdegree}_{\text{gen}}(X)$ is the sum of the degrees of the polar classes $\delta_i(X)$. Here, $\delta_i(X)$ denotes the degree of the *polar class* of X in dimension i , as in [13]. Moreover, if \mathcal{L}^s is a generic linear subspace of codimension s in \mathbb{P}^N , then $\delta_i(X \cap \mathcal{L}^s) = \delta_{i+s}(X)$ by [6, Corollary 6.4]. We call *sth sectional ED degree* of X the number $\text{EDdegree}_{\text{gen}}(X \cap \mathcal{L}^s)$. We denote by X^* the dual variety of X , as in [6, section 5], and already seen in the proof of Proposition 2.2.

COROLLARY 3.2. *The sth sectional ED degree of X is expressed in terms of polar classes as*

$$(3.2) \quad \text{EDdegree}_{\text{gen}}(X \cap \mathcal{L}^s) = \sum_{\ell \geq s} \delta_\ell(X).$$

If $s \leq \text{codim}(X^*) - 1$, then X and $X \cap \mathcal{L}^s$ have the same generic ED degree.

Proof. This follows from results in sections 5 and 6 in [6]. In order to compute $\text{EDdegree}(X \cap \mathcal{L}^s)$ we have to sum $\delta_i(X)$ for $i \geq s$. However, it is known that $\delta_i(X) = 0$ if $i \leq \text{codim}(X^*) - 2$. \square

A special role in [6] is played by the *isotropic quadric* $Q = V(x_0^2 + x_1^2 + \dots + x_N^2)$ in \mathbb{P}^N . If X is smooth and transversal to Q , then [6, Theorem 5.8] gives an explicit formula for the ED degree in terms of Chern classes of X $c_i(X)$. A thorough treatment of Chern classes can be found in [10]; the reader interested in the applications in this paper can be referred to the basics provided in [6]. By combining [6, Theorem 5.8] with Corollary 3.2, we obtain the following.

THEOREM 3.3. *Let $X \subset \mathbb{P}^N$ be a smooth projective variety of dimension M and assume that X is transversal to the isotropic quadric Q . Then the sth sectional ED degree of X equals*

$$\text{EDdegree}_{\text{gen}}(X \cap \mathcal{L}^s) = \sum_{\ell=s}^M \sum_{k=\ell}^M (-1)^{M-k} \binom{k+1}{\ell+1} \text{deg}(c_{M-k}(X)).$$

Proof. The inner sum is the polar class $\delta_i(X)$; see the proof of [6, Theorem 5.8]. \square

We now apply Theorem 3.3 to the situation when $M = m+n-2$, $N = mn-1$, and $X = \mathbb{P}^{m-1} \times \mathbb{P}^{n-1}$ is the Segre variety of $m \times n$ matrices of rank 1 in \mathbb{P}^N . The Chern polynomial of the tangent bundle of X in the Chow ring $A^*(X) = \mathbb{Z}[s, t]/\langle s^m, t^n \rangle$ equals $(1+s)^m(1+t)^n$. By [13, p. 150], this implies

$$(3.3) \quad \delta_\ell(X) = \sum_{k=\ell}^{m+n-2} (-1)^{m+n-k} \binom{k+1}{\ell+1} V_k,$$

where $V_k = \text{deg}(c_{M-k}(X))$ is the coefficient of $s^{m-1}t^{n-1}$ in the expansion of $(1+s)^m(1+t)^n(s+t)^k$. Toric geometers may view V_k as the sum of the normalized volumes of all k -dimensional faces of the polytope $\Delta_{m-1} \times \Delta_{n-1}$; see [6, Corollary 5.11].

The following result explains the ED degrees in the third column in Table 1, and it allows us to determine this column for any desired value of m , n , and s .

THEOREM 3.4. *Let $m \leq n$ and \mathcal{L} be a generic linear subspace of codimension s in $\mathbb{R}^{m \times n}$. For matrices of rank 1 or corank 1, the generic ED degree is given by*

$$(3.4) \quad \begin{aligned} \text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq 1}) &= \delta_s(X) + \delta_{s+1}(X) + \dots + \delta_{m+n-2}(X), \\ \text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq m-1}) &= \delta_0(X) + \delta_1(X) + \dots + \delta_{mn-2-s}(X), \end{aligned}$$

where $\delta_\ell(X)$ may be computed from (3.3).

Proof. The dual in \mathbb{P}^{mn-1} to the Segre variety $X = \mathbb{P}^{m-1} \times \mathbb{P}^{n-1}$ is the variety X^* of matrices of rank $\leq m-1$. By [13, Theorem 2.3], we have $\delta_\ell(X) = \delta_{mn-2-\ell}(X^*)$ for all ℓ . With this duality of polar classes, the result follows from Corollary 3.2 and [6, Theorem 5.4]. \square

Example 3. Fix $m=n=3$. For matrices of rank 1, formulas (3.3) and (3.4) give

$s = \text{codim}(\mathcal{L})$	0	1	2	3	4	5	6	7
V_s	9	18	24	18	6	0	0	0
$\delta_s(X)$	3	6	12	12	6	0	0	0
$\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq 1})$	39	36	30	18	6	0	0	0

Duality for polar classes yields the formulas for 3×3 -matrices of rank $r = 2$ in \mathcal{L} :

$s = \text{codim}(\mathcal{L})$	0	1	2	3	4	5	6	7
$\delta_s(X^*) = \delta_{7-s}(X)$	0	0	0	6	12	12	6	3
$\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq 2})$	39	39	39	39	33	21	9	3

This is our theoretical derivation of the third column in Table 1 for $n = 3$ and generic Λ . \square

Writing down closed formulas for intermediate values of r is more difficult: it involves some Schubert calculus. However, $\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq r})$ can be conveniently computed with the following script in Macaulay2 [12]. It is a slight generalization of that in [6, Example 7.10]:

```
loadPackage "Schubert2"
ED=(m,n,r,s)->
(G = flagBundle({r,m-r}); (S,Q) = G.Bundles;
X=projectiveBundle (S^n); (sx,qx)=X.Bundles;
d=dim X; T=tangentBundle X;
sum(toList(s..m*n-2),i->sum(toList(i..d),j->(-1)^(d-j)*
    binomial(j+1,i+1)*integral(chern(d-j,T)*(chern(1,dual(sx)))^(j))))
```

The function $\text{ED}(m,n,r,s)$ computes the ED degree of the variety of $m \times n$ matrices of rank $\leq r$, in general coordinates, cut with a generic linear space of codimension s in \mathbb{P}^{mn-1} . For $s = 0$, this is precisely the function displayed in [6, Example 7.10].

Example 4. The bold face ED degrees in Table 3 were computed for unit weights $\Lambda = \mathbf{1}$. To find the analogous numbers for generic weights Λ , we run our Macaulay2 code as follows:

```
apply(12,s->ED(4,4,2,s))
{1350, 1350, 1350, 1350, 1330, 1250, 1074, 818, 532, 276, 100, 20}
apply(12,s->ED(3,4,2,s))
{83, 83, 83, 83, 83, 83, 73, 49, 22, 6, 0, 0}
apply(12,s->ED(3,5,2,s))
{143, 143, 143, 143, 143, 143, 143, 143, 128, 88, 40, 10}
```

At this point, we wish to reiterate the main thesis of this paper, namely, that knowing the ED degree ahead of time is useful for practitioners who seek to find and certify the global minimum in the optimization problem (1.3), and to bound the number of local minima. The following example illustrates this for one of the numbers 83 in the output in Example 4.

Example 5. We here solve the generic weighted structured low-rank approximation problem over the reals with parameters $m = 3$, $n = 4$, $r = 2$, and $s = 2$. Consider

the instance

$$U = \begin{bmatrix} -9 & 4 & 9 & -10 \\ 10 & 6 & 1 & -9 \\ 10 & 5 & 7 & 6 \end{bmatrix}, \quad \Lambda = \begin{bmatrix} 8 & 6 & 8 & 2 \\ 1 & 8 & 7 & 9 \\ 7 & 2 & 4 & 6 \end{bmatrix},$$

$$L_1(X) = -10x_{11} + 4x_{12} + 6x_{13} + 8x_{14} + 4x_{21} - 9x_{22} + x_{23} - 10x_{31} - 10x_{32} - 8x_{33} + 2x_{34} - 1,$$

$$L_2(X) = 2x_{11} + 7x_{12} + 3x_{13} - 7x_{14} - 4x_{21} - 6x_{22} - 7x_{23} + 5x_{24} + 8x_{31} + 2x_{33} + 3x_{34} - 1.$$

We wish to find the matrix X of rank at most 2 that satisfies the affine constraints $L_1(X) = L_2(X) = 0$ and is nearest to U . Using Gröbner bases computations and real isolation techniques via the `Maple` packages `FGb` and `fgbrs`, we find that the weighted distance function has 83 complex critical points. This matches the theoretical value $\text{ED}(3, 4, 2, 2) = 83$ provided in Example 4, so that we are guaranteed that there are no further critical points. Among them, seven are real, and we obtain certified numerical approximations of their values:

$$\begin{bmatrix} 0.764 & -1.457 & 2.436 & 1.870 \\ 0.753 & -0.0154 & 0.030 & -7.437 \\ 2.020 & -4.371 & 7.308 & 8.330 \end{bmatrix}, \begin{bmatrix} -8.0341 & 4.127 & 9.055 & 5.364 \\ 16.936 & 2.930 & -1.330 & -4.220 \\ 9.429 & 7.525 & 8.258 & 1.242 \end{bmatrix},$$

$$\begin{bmatrix} -8.215 & 5.033 & 9.965 & 1.647 \\ 16.848 & 4.259 & 0.423 & -3.669 \\ 9.070 & 6.218 & 5.842 & -2.054 \end{bmatrix}, \begin{bmatrix} -8.586 & -1.743 & 1.591 & 2.436 \\ 11.191 & 2.985 & -4.232 & -7.159 \\ 10.351 & 0.292 & 3.567 & 7.185 \end{bmatrix},$$

$$\begin{bmatrix} -4.853 & 4.081 & 6.301 & -6.349 \\ -6.067 & 5.029 & 8.600 & -8.251 \\ 2.616 & -2.455 & -0.878 & 2.327 \end{bmatrix}, \begin{bmatrix} -2.308 & -4.584 & 3.566 & -5.484 \\ -0.205 & -2.210 & 0.668 & -3.178 \\ -2.276 & 0.983 & 2.444 & 2.810 \end{bmatrix},$$

$$\begin{bmatrix} -9.664 & 2.805 & 7.113 & -10.754 \\ 14.942 & 6.520 & 3.149 & -8.783 \\ 8.344 & 0.615 & -2.185 & 2.177 \end{bmatrix}.$$

The last matrix is the closest critical point on the manifold of rank 2 matrices satisfying $L_1 = L_2 = 0$. This computation takes 1002 seconds, and the most time-consuming step is the computation of the Gröbner basis. In order to certify that the global minimum is among these matrices, we also solve the same low-rank approximation problem for rank 1 matrices. Using the same method, this provides us with 11 rank 1 matrices with real entries in 79 seconds. None of them is closer to U than the best rank 2 approximation. Consequently, the global minimum of the weighted distance is reached at the last matrix in the above list.

For comparison purposes, with the same constraints L_1, L_2 and same data matrix U but by taking the Frobenius distance (i.e., Λ is the unit matrix), the number of complex critical points is 43. Five of them are real. Here, it takes only 27 seconds to find the global minimizer. These computations have been performed on an `Intel Xeon E7540/2.00GHz`. \square

In Table 1 and Example 4 we observed that the sectional ED degree for generic Λ does not depend on $s = \text{codim}(\mathcal{L})$, provided s is small. The following corollary explains this.

COROLLARY 3.5. *For a generic linear subspace \mathcal{L} of codimension $s < r(r+n-m)$,*

$$\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq r}) = \text{EDdegree}_{\text{gen}}(\mathbb{C}_{\leq r}^{m \times n}).$$

Proof. Let X be the variety of matrices of rank $\leq r$. Its dual X^* is the variety of matrices of rank $\leq m - r$ and has codimension $\text{codim}(X^*) = (r + n - m)r$. This implies $\delta_\ell(X) = 0$ for $\ell < (r + n - m)r - 1$. The assertion follows from Corollary 3.2. \square

Corollary 3.5 can be stated informally like this: in the setting of generic weights and generic linear spaces of matrices with sufficiently high dimension, the algebraic complexity of structured low-rank approximation agrees with that of ordinary low-rank approximation.

Shifting gears, we now consider the case of unit weights $\Lambda = \mathbf{1}$. Thus, we fix $Q = V(\sum x_{ij}^2)$ as the isotropic quadric in \mathbb{P}^{mn-1} . Let $X = \mathbb{P}^{m-1} \times \mathbb{P}^{n-1}$ denote the Segre variety of $m \times n$ matrices of rank 1 in \mathbb{P}^{mn-1} , and let $Z = \text{Sing}(X \cap Q)$ denote the nontransversal locus of the intersection of X with Q . The dual variety X^* consists of all matrices of rank $\leq m - 1$ in \mathbb{P}^{mn-1} . We conjecture that the following formula (put $m = n$) holds for the gap between the third and the first column of Table 1 (or between the fourth and the second, as well):

$$(3.5) \quad \text{EDdegree}_{\text{gen}}(X^* \cap \mathcal{L}^s) - \text{EDdegree}_1(X^* \cap \mathcal{L}^s) = \text{EDdegree}_{\text{gen}}(Z \cap \mathcal{L}^s).$$

To compute the right-hand side, and to test this conjecture, we use the following.

LEMMA 3.6. *The locus where Q meets $X = \mathbb{P}^{m-1} \times \mathbb{P}^{n-1}$ nontransversally in \mathbb{P}^{mn-1} is the product $Z = Q_{m-2} \times Q_{n-2}$, where Q_{i-2} denotes a general quadratic hypersurface in \mathbb{P}^{i-1} .*

Proof. The Segre variety X meets Q in the union of two irreducible components, $\mathbb{P}^{m-1} \times Q_{n-2}$ and $Q_{m-2} \times \mathbb{P}^{n-1}$. The nontransversality locus is the intersection of these components. \square

Example 6. Let $m = n = 2$, so X and X^* represent 3×3 -matrices of rank 1 and rank ≤ 2 , respectively. Here $Z = Q_1 \times Q_1$ corresponds to the Segre quadric $\mathbb{P}^1 \times \mathbb{P}^1$, embedded in \mathbb{P}^8 with the line bundle $\mathcal{O}(2, 2)$. This is a toric surface whose polygon P is twice a regular square. The facial volumes as in [6, Corollary 5.1] are $V_0 = 4$, $V_1 = 8$, and $V_2 = 8$, and hence

$$\delta_0(Z) = 4 - 2 \cdot 8 + 3 \cdot 8 = 12, \quad \delta_1(Z) = -8 + 3 \cdot 8 = 16, \quad \delta_2(Z) = 8.$$

We fill this into a table and, using Corollary 3.2, we compute the sectional ED degree:

s	0	1	2	3	4	5	6	7
$\delta_s(Z)$	12	16	8	0	0	0	0	0
$\text{EDdegree}_{\text{gen}}(Z \cap \mathcal{L}^s)$	36	24	8	0	0	0	0	0
$\text{EDdegree}_{\text{gen}}(X^* \cap \mathcal{L}^s)$	39	39	39	39	33	21	9	3
$\text{EDdegree}_1(X^* \cap \mathcal{L}^s)$	3	15	31	39	33	21	9	3

The last two lines are taken from Table 1, and they confirm the formula (3.5). \square

Combining Lemma 3.6, Corollary 3.2, and the proof of [6, Theorem 5.8], and abbreviating $W_j = \text{deg}(c_{m+n-4-j}(Q_{m-2} \times Q_{n-2}))$, the right-hand side of (3.5) can be expressed as

$$(3.6) \quad \sum_{i=s}^{m+n-4} \sum_{j=i}^{m+n-4} (-1)^{m+n-4-j} \binom{j+1}{i+1} W_j.$$

Moreover, W_j is equal to the coefficient of $t^{m-2}s^{n-2}$ in the rational generating function

$$4 \frac{(1+t)^m(1+s)^n}{(1+2t)(1+2s)} (t+s)^j.$$

This computation allows us to extend Table 1 to any desired value of m , n , and s .

Changing topics, we now consider the case when \mathcal{L} is the space of *Hankel matrices*. The computation of low-rank approximation of Hankel matrices will be our topic in section 4, where we focus on algebraic geometry and formulas for generic ED degree.

Set $d = p+q-2$ and let $X_{d,r}$ denote the variety of $p \times q$ Hankel matrices of rank $\leq r$. See (4.1) for examples. This variety lives in the projective space $\mathbb{P}^d = \mathbb{P}(S^d \mathbb{C}^2)$, whose points represent binary forms of degree d . Thus $X_{d,1}$ is the rational normal curve of degree d , and $X_{d,r}$ is the r th secant variety of this curve. We have $\dim(X_{d,r}) = 2r - 1$ for $r + 1 \leq \min(p, q)$.

THEOREM 3.7. *Let $d = p + q - 2$ and $r + 1 \leq \min(p, q)$. The generic ED degree of the variety $X_{d,r}$ of $p \times q$ Hankel matrices of rank $\leq r$ in \mathbb{P}^d equals*

$$(3.7) \quad \text{EDdegree}_{\text{gen}}(X_{d,r}) = \sum_{i=0}^r \binom{d+1-r}{i} \binom{d-r-i}{r-i} 2^{r-i}.$$

Proof. The sum in (3.7) is the coefficient of z^r in the generating function

$$(3.8) \quad \frac{(1+z)^{d+1-r}}{(1-2z)^{d-2r+1}}.$$

The conormal variety of $X_{d,r}$ is the closure $\mathcal{N}_{X_{d,r}}$ of the set

$$\{(f, g) \mid \text{rank}(f) = r \text{ and } g \text{ is tangent to } X_{d,r} \text{ at } f\} \subset \mathbb{P}(S^d \mathbb{C}^2) \times \mathbb{P}(S^d(\mathbb{C}^2)^*).$$

The homology class of $\mathcal{N}_{X_{d,r}}$ is given by a binary form. We will show that the sum $\sum_i \delta_i(X_{d,r})$ of its coefficients is the asserted coefficient of (3.8). By [6, (5.3)], this proves the claim.

Let p_1, p_2 be the two projections. The images of the conormal variety $\mathcal{N}_{X_{d,r}}$ are

$$p_1(\mathcal{N}_{X_{d,r}}) = X_{d,r} \quad \text{and} \quad p_2(\mathcal{N}_{X_{d,r}}) = X_{d,r}^*.$$

We desingularize $X_{d,r}$ by considering $\text{Sym}^r(\mathbb{P}^1) \simeq \mathbb{P}^r$. The desingularization map is given by the scheme-theoretic intersection of the rational normal curve of degree r with a hyperplane. A point in \mathbb{P}^r , identified with a hyperplane, gives r points on $X_{d,1} \simeq \mathbb{P}^1$. Their linear span in \mathbb{P}^d defines a rank r bundle on \mathbb{P}^r , known as the *Schwarzenberger bundle* [7, section 6]. This is the kernel of the bundle map $\mathcal{O}^{d+1} \rightarrow \mathcal{O}(1)^{d+1-r}$. In the same way, we desingularize the conormal variety $\mathcal{N}_{X_{d,r}}$ by the fiber product over \mathbb{P}^r of the projectivization of the Schwarzenberger bundle $E_{d,r} = \text{kernel}(\mathcal{O}^{d+1} \rightarrow \mathcal{O}(1)^{d+1-r})$ and of the projective bundle of $\mathcal{O}(2)^{d-2r+1}$. Exactly as in the proof of [3, Proposition 4.1], the degrees of the polar classes of $X_{d,r}$ are

$$\delta_{r+i-1}(X_{d,r}) = \int_{\mathbb{P}^r} s_i(E_{d,r}) s_{r-i}(\mathcal{O}(2)^{d-2r+1}).$$

The total Segre class of $E_{d,r}$ is $(1+z)^{d+1-r}$. The total Segre class of $\mathcal{O}(2)^{d-2r+1}$ is $\frac{1}{(1-2z)^{d-2r+1}}$. By multiplying them we obtain the degree sum of the polar classes, thus proving (3.8). \square

COROLLARY 3.8. *The generic ED degree of the hypersurface $X_{2r,r}$ defined by the Hankel determinant of format $(r+1) \times (r+1)$ is equal to*

$$(3.9) \quad \frac{3^{r+1} - 1}{2} = \text{the coefficient of } z^r \text{ in } \frac{(1+z)^{r+1}}{1-2z}.$$

This corollary means that the ED degree of the $(r+1) \times (r+1)$ Hankel determinant agrees with the ED degree of the general symmetric $(r+1) \times (r+1)$ determinant. By ED duality [6, Theorem 5.2], this also the ED degree of the second Veronese embedding of \mathbb{P}^r ; see [6, Example 5.6]. If we consider Hankel matrices of fixed rank r , then we obtain polynomiality:

COROLLARY 3.9. *For fixed r , the generic ED degree of $X_{d,r}$ is a polynomial of degree r in d .*

For example, we find the following explicit polynomials when the rank r is small:

$$\begin{aligned} \text{EDdegree}_{\text{gen}}(X_{d,1}) &= 3d - 2, \\ \text{EDdegree}_{\text{gen}}(X_{d,2}) &= (9d^2 - 39d + 38)/2, \\ \text{EDdegree}_{\text{gen}}(X_{d,3}) &= (9d^3 - 99d^2 + 348d - 388)/2, \\ \text{EDdegree}_{\text{gen}}(X_{d,4}) &= (27d^4 - 558d^3 + 4221d^2 - 13818d + 16472)/8. \end{aligned}$$

The values of these polynomials are the entries in the left columns in Table 4 below.

4. Hankel and Sylvester matrices. In this section we study the weighted low-rank approximation problem for matrices with a special structure that is given by equating some matrix entries and setting others to zero. One such family consists of the Hurwitz matrices in [6, Theorem 3.6]. We here discuss Hankel matrices, then catalecticants, and finally Sylvester matrices. The corresponding applications are low-rank approximation of symmetric tensors and approximate greatest common divisors.

The *Hankel matrix* $H[p, q]$ of format $p \times q$ has the entry x_{i+j-1} in row i and column j . So, the total number of unknowns is $n = p + q - 1$. We are most interested in the case when this matrix is square or almost square. The *Hankel matrix of order n* is $H[(n+1)/2, (n+1)/2]$ if n is odd, and it is $H[(n/2, (n+2)/2]$ if n is even. We denote this matrix by H_n . For instance,

$$(4.1) \quad H_5 = \begin{bmatrix} x_1 & x_2 & x_3 \\ x_2 & x_3 & x_4 \\ x_3 & x_4 & x_5 \end{bmatrix} \quad \text{and} \quad H_6 = \begin{bmatrix} x_1 & x_2 & x_3 & x_4 \\ x_2 & x_3 & x_4 & x_5 \\ x_3 & x_4 & x_5 & x_6 \end{bmatrix}.$$

For approximations by low-rank Hankel matrices, we consider three natural weights:

- the matrix Ω_n has entry $1/\min(i+j-1, n-i-j+2)$ in row i and column j ;
- the matrix $\mathbf{1}_n$ has all entries equal to 1;
- the matrix Θ_n has $\binom{n-1}{i+j-2}/\min(i+j-1, n-i-j+2)$ in row i and column j .

We encountered these matrices for $n = 5$ in Example 1. For $n = 6$ we have

$$\Omega_6 = \begin{bmatrix} 1 & 1/2 & 1/3 & 1/3 \\ 1/2 & 1/3 & 1/3 & 1/2 \\ 1/3 & 1/3 & 1/2 & 1 \end{bmatrix} \quad \text{and} \quad \Theta_6 = \begin{bmatrix} 1 & 5/2 & 10/3 & 10/3 \\ 5/2 & 10/3 & 10/3 & 5/2 \\ 10/3 & 10/3 & 5/2 & 1 \end{bmatrix}.$$

The weights Ω_n represent the usual Euclidean distance in \mathbb{R}^n , the unit weights $\mathbf{1}_n$ give the Frobenius distance in the ambient matrix space, and the weights Θ_n give the natural metric in the space of symmetric $2 \times 2 \times \dots \times 2$ -tensors. Such a tensor corresponds to a binary form

$$F(s, t) = \sum_{i=1}^n \binom{n-1}{i-1} \cdot x_i \cdot s^{n-i} \cdot t^{i-1}.$$

The Hankel matrix H_n has rank 1 if and only if $F(s, t)$ is the $(n-1)$ st power of a linear form. More generally, if $F(s, t)$ is the sum of r powers of linear forms, then H_n has rank $\leq r$. As we saw in section 3, this locus corresponds to the r th secant variety of

TABLE 4
 Weighted ED degrees for Hankel matrices of order n and rank r .

$\Lambda = \Omega_n$				
$n \setminus r$	1	2	3	4
3	4			
4	7			
5	10	13		
6	13	34		
7	16	64	40	
8	19	103	142	
9	22	151	334	121

$\Lambda = \mathbf{1}_n$				
$n \setminus r$	1	2	3	4
3	2			
4	7			
5	6	9		
6	13	34		
7	10	38	34	
8	19	103	142	
9	14	103	246	113

$\Lambda = \Theta_n$				
$n \setminus r$	1	2	3	4
3	2			
4	3			
5	4	7		
6	5	16		
7	6	28	20	
8	7	43	62	
9	8	61	134	53

the rational normal curve in \mathbb{P}^{n-1} . Various ED degrees for our three weight matrices are displayed in Table 4.

The entries in the leftmost chart in Table 4 come from Theorem 3.7. Indeed, the variety of Hankel matrices H_n of rank $\leq r$ is precisely the secant variety $X_{n-1,r}$ we discussed in section 3. The weight matrix $\Lambda = \Omega_n$ exhibits the generic ED degree for that variety. The columns on the left of Table 4 are the values of the polynomials in Corollary 3.9, and the diagonal entries 4, 13, 40, 121, . . . are given by Corollary 3.8.

All ED degrees in Table 4 were verified using Gröbner basis computations over $\text{GF}(65521)$ using the `maple` package `FGb` [8]. The running times are closely tied to the value of the ED degrees, and they are similar to those reported in Table 3. Gröbner bases over \mathbb{Q} can also be computed fairly easily whenever the ED degree is below 100, and for those cases we can locate all real critical points using `fgbrs`. However, for larger instances, exact symbolic solving over \mathbb{Q} becomes a considerable challenge due to the growth in coefficient size.

Hankel matrices of rank r correspond to symmetric $2 \times 2 \times \dots \times 2$ -tensors of tensor rank r , and these can be represented by binary forms that are sums of r powers of linear forms. That is the point of the geometric discussion in section 3. This interpretation extends to symmetric tensors of arbitrary format, with the rational normal curve replaced with the Veronese variety. For a general study of low-rank approximation of symmetric tensors, see Friedland and Stawiska [9]. In general, there is no straightforward representation of low rank tensors by low rank matrices with special structure. However, there are some exceptions, notably for rank $r = 2$ tensors, by the results of Raicu [20] and others in the recent tensor literature. We refer to Landsberg’s book [16], especially Chapters 3, 7, and 10. The resulting generalized Hankel matrices are known as *catalecticants* in the commutative algebra literature, or as *moment matrices* in the optimization literature. We now present a case study that arose from a particular application in biomedical imaging.

We consider the following catalecticant matrix of format 6×6 :

$$X = \begin{bmatrix} x_{400} & x_{310} & x_{301} & x_{220} & x_{211} & x_{202} \\ x_{310} & x_{220} & x_{211} & x_{130} & x_{121} & x_{112} \\ x_{301} & x_{211} & x_{202} & x_{121} & x_{112} & x_{103} \\ x_{220} & x_{130} & x_{121} & x_{040} & x_{031} & x_{022} \\ x_{211} & x_{121} & x_{112} & x_{031} & x_{022} & x_{013} \\ x_{202} & x_{112} & x_{103} & x_{022} & x_{013} & x_{004} \end{bmatrix}.$$

The fifteen unknown entries are the coefficients of a ternary quartic

$$\begin{aligned} F(s, t, u) = & x_{400}s^4 + x_{040}t^4 + x_{004}u^4 + 6x_{220}s^2t^2 + 6x_{202}s^2u^2 + 6x_{022}t^2u^2 \\ & + 4x_{310}s^3t + 4x_{301}s^3u + 4x_{130}st^3 + 4x_{031}t^3u + 4x_{103}su^3 \\ & + 4x_{013}tu^3 + 12x_{211}s^2tu + 12x_{121}st^2u + 12x_{112}stu^2. \end{aligned}$$

The table (x_{ijk}) can be regarded as a symmetric tensor of format $3 \times 3 \times 3 \times 3$. The coefficients in $F(s, t, u)$ indicate the multiplicity with which the 15 unknowns occur among the $3^4 = 81$ coordinates of that tensor. To model the invariant metric in the tensor space $\mathbb{R}^{3 \times 3 \times 3 \times 3}$ in our matrix representation, we use the weight matrix

$$\Theta = \begin{bmatrix} 1 & 2 & 2 & 2 & 3 & 2 \\ 2 & 2 & 3 & 2 & 3 & 3 \\ 2 & 3 & 2 & 3 & 3 & 2 \\ 2 & 2 & 3 & 1 & 2 & 2 \\ 3 & 3 & 3 & 2 & 2 & 2 \\ 2 & 3 & 2 & 2 & 2 & 1 \end{bmatrix}.$$

The problem is to approximate a given catalecticant matrix $U = (u_{ijk})$ by a rank 2 matrix with respect to Θ . The expected number of critical points is as follows.

PROPOSITION 4.1. *Let \mathcal{L} be the 15-dimensional subspace of catalecticants X in $\mathbb{R}^{6 \times 6}$. Then*

$$\text{EDdegree}_{\Theta}(\mathcal{L}_{\leq 2}) = 195 \quad \text{and} \quad \text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq 2}) = 1813.$$

The proof is a computation as explained below. We first discuss an application.

Example 7. We consider the following symmetric $3 \times 3 \times 3 \times 3$ -tensor:

u_{400}	0.1023	u_{040}	0.0197	u_{004}	0.1869
u_{220}	0.0039	u_{202}	0.0407	u_{022}	-0.00017418
u_{310}	-0.002	u_{301}	0.0581	u_{130}	0.0107
u_{103}	0.0196	u_{031}	0.0029	u_{013}	-0.0021
u_{211}	-0.00032569	u_{121}	-0.0012	u_{112}	-0.0011

This tensor was given to us by Thomas Schultz, who heads the Visualization and Medical Image Analysis Group at the University of Bonn. It represents a fiber distribution function, estimated from diffusion magnetic resonance imaging. See [23] for more information. \square

We present an algebraic formulation of our problem, which was found to be suitable for symbolic computation. Introducing six unknowns a, b, c, d, e, f , we parametrize the six-dimensional variety of symmetric $3 \times 3 \times 3 \times 3$ -tensors of rank 2 by the ternary quartics

$$\tilde{F}(s, t, u) = a \cdot (s + bt + cu)^4 + d \cdot (s + et + fu)^4.$$

Just like in the discussion in Remark 1 and after Proposition 2.3, the image of this parametrization is a dense open subset of the symmetric $3 \times 3 \times 3 \times 3$ -tensors of rank 2. Covering all rank 2 tensors can be achieved with three parametrizations as above.

Written out explicitly, this parametrization takes the form

$$\begin{array}{lll} x_{400} = a + d, & x_{040} = ab^4 + de^4, & x_{004} = ac^4 + df^4, \\ x_{220} = ab^2 + de^2, & x_{202} = ac^2 + df^2, & x_{022} = ab^2c^2 + de^2f^2, \\ x_{310} = ab + de, & x_{301} = ac + df, & x_{130} = ab^3 + de^3, \\ x_{103} = ac^3 + df^3, & x_{031} = ab^3c + de^3f, & x_{013} = abc^3 + def^3, \\ x_{211} = abc + def, & x_{121} = ab^2c + de^2f, & x_{112} = abc^2 + def^2. \end{array}$$

Note that our parametrization is 2 to 1: every rank 2 catalecticant X has two preimages, which are related by swapping the vectors (a, b, c) and (d, e, f) . The fiber

jumps in dimension over the singular locus, which consists of matrices X of rank 1. Their preimage in parameter space is given by the ideal $\langle ad \rangle \cap \langle b - e, c - f \rangle$. The chosen weight matrix Θ now specifies the following unconstrained optimization problem. We seek to find the minimum in \mathbb{R}^6 of

$$\begin{aligned} G(a, b, c, d, e, f) &= (u_{400} - a - d)^2 + (u_{040} - ab^4 - de^4)^2 + (u_{004} - ac^4 - df^4)^2 \\ &+ 6(u_{220} - ab^2 - de^2)^2 + 6(u_{202} - ac^2 - df^2)^2 + 6(u_{022} - ab^2c^2 - de^2f^2)^2 \\ &+ 4(u_{310} - ab - de)^2 + 4(u_{301} - ac - df)^2 + 4(u_{130} - ab^3 - de^3)^2 \\ &+ 4(u_{103} - ac^3 - df^3)^2 + 4(u_{031} - ab^3c - de^3f)^2 + 4(u_{013} - abc^3 - def^3)^2 \\ &+ 12(u_{112} - abc^2 - def^2)^2 + 12(u_{211} - abc - def)^2 + 12(u_{121} - ab^2c - de^2f)^2. \end{aligned}$$

The set of complex critical points is the zero locus of the ideal

$$I = \left\langle \frac{\partial G}{\partial a}, \frac{\partial G}{\partial b}, \frac{\partial G}{\partial c}, \frac{\partial G}{\partial d}, \frac{\partial G}{\partial e}, \frac{\partial G}{\partial f} \right\rangle : (\langle ad \rangle \cap \langle b - e, c - f \rangle)^\infty.$$

For applications, we are interested in the real points in this variety.

Computational proof of Proposition 4.1. As argued in [6, section 2], the ideal I is radical and zero-dimensional when the u_{ijk} are generic rational numbers. The number of solutions is the degree of I , and we found this to be $370 = 2 \cdot 195$. This is twice the ED degree of $\mathcal{L}_{\leq 2}$ with respect to $\Lambda = \Theta$. For this computation we used the `FGb` library in `maple`. We used Gröbner bases over the finite field $\text{GF}(65521)$ to avoid the swelling of rational coefficients, the data u_{ijk} are chosen uniformly at random in this field, and we saturate only by $\langle ad(b - e) \rangle$. The computation took 90 seconds and returned 390 critical points of G . Performing the same computation with the coefficients 1, 6, 4, 12 in $G(a, b, c, d, e, f)$ replaced with random field elements, we find $3626 = 2 \cdot 1813$ critical points, and hence $\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq 2}) = 1813$. \square

Example 8. We return to the particular data set in Example 7. Using the above parametrization, the best rank 2 approximation can be obtained by solving a polynomial system. This can be achieved by using symbolic or numerical methods.

A numerical computation conducted by Jose Rodriguez with the software `Bertini` indicates that, for Thomas Schultz's data, precisely 9 of the 195 critical points are real. These correspond to 2 local minima and 7 saddle points of the Euclidean distance function. The precomputation with generic data took 2 hours on 40 AMD Opteron 6276/2.3Ghz cores. Then the computation with the numerical data in Example 7 was achieved in 1 minute.

These results were also computed by symbolic methods: a Gröbner basis computation conducted by Jean-Charles Faugère and Mohab Safey El Din with the software `FGb` returned an algebraic parametrization of the 195 complex critical points by the roots of a univariate polynomial of degree 195. This polynomial has 9 real roots. Two of them correspond to the two local minima. The average size of the integer coefficients of this univariate polynomial is 11000 digits. For this computation, the above formulation as an unconstrained optimization problem was used. It took 11 minutes on a 2.6GHz IntelCore i7. In general, for symbolic methods unconstrained formulations seem to be better than the general implicit formulation in Proposition 2.3. See the comparisons of timings in Table 3. However, most instances of (1.3) do not admit an unconstrained formulation, because $\mathcal{L}_{\leq r}$ is usually not unirational. \square

Our last topic in this section is the study of Sylvester matrices. We consider two arbitrary polynomials F and G in one variable t . Suppose their degrees are m and n

with $m \leq n$, so

$$F(t) = \sum_{i=0}^m a_i t^i \quad \text{and} \quad G(t) = \sum_{j=0}^n b_j t^j.$$

Fix k with $1 \leq k \leq m$. The k th Sylvester matrix of the pair (F, G) equals

$$\text{Syl}_k(F, G) = \begin{bmatrix} a_0 & 0 & \cdots & 0 & b_0 & 0 & \cdots & 0 \\ \vdots & a_0 & \ddots & \vdots & \vdots & b_0 & \ddots & \vdots \\ a_m & \vdots & \ddots & 0 & b_n & \vdots & \ddots & 0 \\ 0 & a_m & \vdots & a_0 & 0 & b_n & \vdots & b_0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_m & 0 & 0 & \cdots & b_n \end{bmatrix}.$$

This matrix has $n + k$ rows and $n - m + 2k$ columns, so it is square for $k = m$, and it has more rows than columns for $k < m$. The maximal minors have size $n - m + 2k$, and they all vanish when $\text{Syl}_k(F, G)$ has a nonzero vector in its kernel. Such a vector corresponds to a polynomial of degree $m - k + 1$ that is a common factor of F and G .

The *approximate gcd problem* in computer algebra [14, 15] aims to approximate a given pair (F, G) by a nearby pair (F^*, G^*) whose Sylvester matrix $\text{Syl}_k(F^*, G^*)$ has linearly dependent columns. Writing \mathcal{L} for the subspace of Sylvester matrices, this is precisely our ED problem for $\mathcal{L}_{\leq n-m+2k-1}$. The following theorem furnishes a formula for $\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq n-m+2k-1})$.

THEOREM 4.2. *For the variety of pairs (F, G) of univariate polynomials of degrees (m, n) with a common factor of degree $m - k + 1$, the generic ED degree equals that of the Segre variety of $(m - k + 2) \times (n - m + 2k)$ -matrices of rank 1. It is given by setting $s = 0$ in (3.4). Using the `Macaulay2` function `ED` in Example 4, we can write this ED degree as*

$$\text{EDdegree}_{\text{gen}}(\mathcal{L}_{\leq n-m+2k-1}) = \text{ED}(m - k + 2, n - m + 2 * k, 1, 0).$$

Proof. A natural desingularization is given by multiplying with the desired common factor:

$$(4.2) \quad \begin{array}{ccc} \mathbb{P}^{m-k+1} \times \mathbb{P}^{n-m+2k-1} & \rightarrow & \mathcal{L}_{\leq n-m+2k-1}, \\ [A(t), (B(t), C(t))] & \mapsto & [A(t)B(t), A(t)C(t)]. \end{array}$$

Here $A(t), B(t), C(t)$ are polynomials of degrees $m - k + 1, k - 1, n - m + k - 1$ respectively. The map (4.2) lifts to a linear projection map from the Segre embedding of $\mathbb{P}^{m-k+1} \times \mathbb{P}^{n-m+2k-1}$. Work of Piene [19, section 4] implies that the degrees of polar loci can be computed on that Segre variety. The ED degree is a sum of degrees of these, by Corollary 3.2. The result follows. \square

For $m = k$, when the Sylvester matrix is square, Theorem 4.2 refers to $2 \times (n + m)$ -matrices of rank 1. Similarly to [6, Example 5.12], their ED degree is $4(m + n) - 2$.

COROLLARY 4.3. *The generic ED degree of the Sylvester determinant Syl_m equals $4(m + n) - 2$.*

We consider three natural choices of weight matrices for the low-rank approximation of Sylvester matrices. As before in Table 4, we write $\Omega_{m,n}$ for the weight matrix that represents the Euclidean distance on \mathbb{R}^{m+n+2} : it is the matrix which has the

TABLE 5
 Weighted ED degrees for Sylvester matrices $\text{Syl}_k(F, G)$.

Λ is generic					$\Lambda = \Omega_{m,n}$					$\Lambda = \Theta_{m,n}$				
$(m, n) \setminus k$	1	2	3	4	$(m, n) \setminus k$	1	2	3	4	$(m, n) \setminus k$	1	2	3	4
(2, 2)	10	14			(2, 2)	2	6			(2, 2)	2	4		
(2, 3)	39	18			(2, 3)	23	18			(2, 3)	19	6		
(2, 4)	83	22			(2, 4)	75	22			(2, 4)	29	8		
(2, 5)	143	26			(2, 5)	119	18			(2, 5)	61	10		
(3, 3)	14	83	22		(3, 3)	2	19	10		(3, 3)	2	19	6	
(3, 4)	83	143	26		(3, 4)	35	95	26		(3, 4)	41	53	8	
(3, 5)	284	219	30		(3, 5)	188	203	26		(3, 5)	106	81	10	
(4, 4)	18	284	219	30	(4, 4)	2	36	59	14	(4, 4)	2	50	45	8
(4, 5)	143	676	311	34	(4, 5)	47	276	215	34	(4, 5)	71	256	101	10

same pattern as Syl_k with a_i and b_j replaced, respectively, by $1/(n - m + k)$ and $1/k$. We also write $\Theta_{m,n}$ for the weight matrix of the rotation invariant quadratic form: a_i is replaced by $1/((n - m + k) \binom{m}{i})$ and b_j is replaced by $1/(k \binom{n}{j})$. In Table 5 we present the ED degrees for these choices of weights. The left table shows the generic behavior predicted by Theorem 4.2. At present, we do not know a general formula for the entries of the two tables on the right side, but we are hopeful that an approach like (3.5) will lead to such formulas. Along the rightmost margins, where the matrix Syl_m is square, the formula seems to be $\text{EDdegree}_\Theta(\mathcal{L}_{\leq n+k-1}) = 2n$.

Acknowledgments. We thank the following colleagues for their help with this project: Jean-Charles Faugère, William Rey, Ragni Piene, Jose Rodriguez, Mohab Safey El Din, Éric Schost, and Thomas Schultze. Giorgio Ottaviani is a member of GNSAGA-INDAM. Pierre-Jean Spaenlehauer and Bernd Sturmfels were hosted by the Max-Planck Institute für Mathematik in Bonn, Germany.

REFERENCES

- [1] D. BATES, J. HAUENSTEIN, A. SOMMESE, AND C. WAMPLER, *Numerically Solving Polynomial Systems with Bertini*, SIAM, Philadelphia, 2013.
- [2] G. BLEKHERMAN, P. PARRILO, AND R. THOMAS, *Semidefinite Optimization and Convex Algebraic Geometry*, MOS-SIAM Ser. Optim. 13, SIAM, Philadelphia, 2013.
- [3] H.-C.G. VON BOTHMER, AND K. RANESTAD, *A general formula for the algebraic degree in semidefinite programming*, Bull. Lond. Math. Soc., 41 (2009), pp. 193–197.
- [4] F. CATANESE, S. HOŞTEN, A. KHETAN, AND B. STURMFELS, *The maximum likelihood degree*, Amer. J. Math., 128 (2006), pp. 671–697.
- [5] M. CHU, R. FUNDERLIC, AND R. PLEMMONS, *Structured low rank approximation*, Linear Algebra Appl., 366 (2003), pp. 157–172.
- [6] J. DRAISMA, E. HOROBEŢ, G. OTTAVIANI, B. STURMFELS, AND R. THOMAS, *The Euclidean distance degree of an algebraic variety*, Found. Comput. Math., to appear.
- [7] I. DOLGACHEV AND M. KAPRANOV, *Arrangement of hyperplanes and vector bundles on \mathbb{P}^n* , Duke Math. J., 71 (1993), pp. 633–664.
- [8] J.-C. FAUGÈRE, *A new efficient algorithm for computing Gröbner bases without reduction to zero (F5)*, in Proceedings of ISSAC 2002, pp. 75–83.
- [9] S. FRIEDLAND AND M. STAWISKA, *Best approximation on semi-algebraic sets and k -border rank approximation of symmetric tensors*, arXiv:1311.1561, 2013.
- [10] W. FULTON, *Intersection Theory*, Springer, Berlin, 1998.
- [11] M. GOLUBITSKY AND V. GUILLEMIN, *Stable Mappings and their Singularities*, Springer, Berlin, 1974.
- [12] D.R. GRAYSON AND M.E. STILLMAN, *Macaulay2, a Software System for Research in Algebraic Geometry*, <http://www.math.uiuc.edu/Macaulay2/Citing/>.
- [13] A. HOLME, *The geometric and numerical properties of duality in projective algebraic geometry*, Manuscripta Math., 61 (1988), pp. 145–162.

- [14] N.K. KARMARKAR AND Y.N. LAKSHMAN, *On approximate GCDs of univariate polynomials*, J. Symbolic Comput., 26 (1998), pp. 653–666.
- [15] E. KALTOFEN, Z. YANG, AND L. ZHI, *Structured low rank approximation of a Sylvester matrix*, in Symbolic-Numeric Computation, Trends in Mathematics, D. Wang and L. Zhi, eds., Birkhäuser, Basel, Switzerland, 2007, pp. 69–83.
- [16] J.M. LANDSBERG, *Tensors: Geometry and Applications*, Grad. Stud. Math. 128, AMS, Providence, RI, 2012.
- [17] J.H. MANTON, R. MAHONY, AND Y. HUA, *The geometry of weighted low-rank approximation*, IEEE Trans. Signal Process., 51 (2003), pp. 500–514.
- [18] I. MARKOVSKI, *Structured low-rank approximation and its applications*, Automatica J. IFAC, 44 (2008), pp. 891–909.
- [19] R. PIENE, *Polar classes of singular varieties*, Ann. Sci. Éc. Norm. Supér. (4), 11 (1978), pp. 247–276.
- [20] C. RAICU, *Secant varieties of Segre-Veronese varieties*, Algebra Number Theory, 6 (2012), pp. 1817–1868.
- [21] K. RANESTAD, *Algebraic degree in semidefinite and polynomial optimization*, in Handbook on Semidefinite, Conic and Polynomial Optimization, J.-B. Lasserre and M. Anjos, eds., Springer, Berlin, 2012, pp. 61–75.
- [22] W. REY, *On weighted low-rank approximation*, arXiv:1302.0360, 2013.
- [23] T. SCHULTZ, A. FUSTER, A. GHOSH, R. DERICHE, L. FLORACK, AND L.-H. LIM, *Higher-order tensors in diffusion imaging*, in Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data, Springer, Berlin, 2013, pp. 129–161.
- [24] N. SREBRO AND T. JAAKKOLA, *Weighted low-rank approximations*, in Proceedings of the International Conference on Machine Learning, 2003, pp. 720–727.