Singular value decomposition of complexes

Danielle A. Brake, Jonathan D. Hauenstein, Frank-Olaf Schreyer,
Andrew J. Sommese, and Michael E. Stillman

May 19, 2018

Abstract

Singular value decompositions of matrices are widely used in numerical linear algebra with many applications. In this paper, we extend the notion of singular value decompositions to finite complexes of real vector spaces. We provide two methods to compute them and present several applications.

1 Introduction

For a matrix $A \in \mathbb{R}^{m \times k}$, a singular value decomposition (SVD) of A is

$$A = U \Sigma V^t$$

where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{k \times k}$ are orthogonal and $\Sigma \in \mathbb{R}^{m \times k}$ is diagonal with nonnegative real numbers on the diagonal. The diagonal entries of Σ, say $\sigma_1 \geq \cdots \geq \sigma_{\min\{m,k\}} \geq 0$ are called the singular values of A and the number of nonzero singular values is equal to the rank of A. Many problems in numerical linear algebra can be solved using a singular value decomposition such as pseudoinversion, least squares solving, and low-rank matrix approximation.

A matrix $A \in \mathbb{R}^{m \times k}$ defines a linear map $A : \mathbb{R}^k \to \mathbb{R}^m$ via $x \mapsto Ax$ denoted

$$\mathbb{R}^m \xleftarrow{\ A\ } \mathbb{R}^k.$$

Hence, matrix multiplication simply corresponds to function composition. For example, if $B \in \mathbb{R}^{\ell \times m}$, then $B \circ A : \mathbb{R}^k \to \mathbb{R}^\ell$ is defined by $x \mapsto BAx$ denoted

$$\mathbb{R}^\ell \xleftarrow{\ B\ } \mathbb{R}^m \xleftarrow{\ A\ } \mathbb{R}^k.$$

If $B \circ A = 0$, then this composition forms a *complex* denoted

$$0 \longleftarrow \mathbb{R}^\ell \xleftarrow{\; B \;} \mathbb{R}^m \xleftarrow{\; A \;} \mathbb{R}^k \longleftarrow 0.$$

In general, a finite complex of finite dimensional $\mathbb{R}$-vector spaces

$$0 \longleftarrow C_0 \xleftarrow{\; A_1 \;} C_1 \xleftarrow{\; A_2 \;} \ldots \xleftarrow{\; A_{n-1} \;} C_{n-1} \xleftarrow{\; A_n \;} C_n \longleftarrow 0$$

consists of vector spaces $C_i \cong \mathbb{R}^{c_i}$ and differentials given by matrices A_i so that $A_i \circ A_{i+1} = 0$. We denote such a complex by $C_\bullet$ and its i^{th} homology group as

$$H_i = H_i(C_\bullet) = \frac{\ker A_i}{\text{image } A_{i+1}}$$

with $h_i = \dim H_i$. Complexes are standard tools that occur in many areas of mathematics including differential equations, e.g., [AFW06, AFW10]. One of the reasons for developing a singular value decomposition of complexes is to compute the dimensions h_i efficiently and robustly via numerical methods when each A_i is only known approximately, say B_i. For example, if we know that $\text{rank } A_i = r_i$, then h_i could easily be computed via

$$c_i = r_i + r_{i+1} + h_i.$$

One option would be to compute the singular value decomposition of each B_i in order to compute the rank r_i of A_i since the singular value decomposition is an excellent rank-revealing numerical method. However, simply decomposing each B_i ignores the important information that the underlying matrices A_i form a complex.

The key point of this paper is that we can utilize information about the complex to provide more specific information that reflects the structure it imposes.

Theorem 1.1 (Singular value decomposition of complexes). *Let $A_1, \ldots, A_n$ be a sequence of matrices $A_i \in \mathbb{R}^{c_{i-1} \times c_i}$ which define a complex $C_\bullet$, i.e. $A_i \circ A_{i+1} = 0$. Let $r_i = \text{rank } A_i$ and $h_i = c_i - (r_i + r_{i+1})$. Then, there exists sequences $U_0, \ldots, U_n$ and $\Sigma_1, \ldots, \Sigma_n$ of orthogonal and diagonal matrices, respectively, such that*

$$U_{i-1}^t \circ A_i \circ U_i = \begin{matrix} & \begin{matrix} r_i & r_{i+1} & h_i \end{matrix} \\ \begin{matrix} r_{i-1} \\ r_i \\ h_{i-1} \end{matrix} & \begin{pmatrix} 0 & 0 & 0 \\ \Sigma_i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \end{matrix} := \overline{\Sigma}_i$$

where all diagonal entries of Σ_i are strictly positive. Moreover, if every $r_i > 0$ and at least one $h_i > 0$, then the orthogonal matrices U_i can be chosen such that $\det U_i = 1$, i.e., each U_i is a special orthogonal matrix.

The diagonal entries of $\Sigma_1, \ldots, \Sigma_n$ are the *singular values of the complex.*

We develop two methods that utilize the complex structure to compute a singular value decomposition of $C_\bullet$. The successive projection method described in Algorithm 3.1 uses the orthogonal projection

$$P_{i-1} \colon C_{i-1} \to \ker A_{i-1}$$

together with the singular value decomposition of $P_{i-1} \circ A_i$. The second method, described in Algorithm 3.3, is based on the *Laplacians*

$$\Delta_i = A_i^t \circ A_i + A_{i+1} \circ A_{i+1}^t.$$

Both methods can be applied to numerical approximations B_i of A_i.

Organization of this paper is as follows. Section 2 proves Theorem 1.1 and collects a number of basic facts along with defining the pseudoinverse of a complex. Section 3 describes the algorithms mentioned above and illustrates them on an example. Section 4 considers projecting an arbitrary sequence of matrices onto a complex. Section 5 provides an application to computing Betti numbers of minimal free resolutions of graded modules over the polynomial ring $\mathbb{Q}[x_0, \ldots, x_n]$ which combines our method with ideas from [EMSS16].

2   Basics

We start with a proof of our main theorem.

Proof of Theorem 1.1. For convenience, we set $A_0 = A_{n+1} = 0$ to compliment $A_1, \ldots, A_n$ that describe the complex. By the homomorphism theorem

$$(\ker A_i)^\perp \cong \mathrm{image}\, A_i.$$

The singular value decomposition for a complex follows by applying singular value decomposition to this isomorphism and extending an orthonormal basis of

these spaces to orthonormal basis of $\mathbb{R}^{c_{i-1}}$ and $\mathbb{R}^{c_i}$. Since image $A_{i+1} \subset \ker A_i$ we have an orthogonal direct sum

$$(\ker A_i)^\perp \oplus \text{image } A_{i+1} \subset \mathbb{R}^{c_i}$$

with

$$H_i := ((\ker A_i)^\perp \oplus \text{image } A_{i+1})^\perp = \ker A_i \cap \text{image } A_{i+1}^\perp \cong \frac{\ker A_i}{\text{image } A_{i+1}}.$$

With respect to these subspaces, we can decompose A_i as

$$
\begin{array}{c}
 \\
(\ker A_{i-1})^\perp \\
\text{image } A_i \\
H_{i-1}
\end{array}
\begin{array}{ccc}
(\ker A_i)^\perp & \text{image } A_{i+1} & H_i \\
\left(\begin{array}{ccc}
0 & 0 & 0 \\
\Sigma_i & 0 & 0 \\
0 & 0 & 0
\end{array}\right).
\end{array}
$$

Indeed, A_i has no component mapping to $(\text{image } A_i)^\perp$, which explains six of the zero blocks, and $\ker A_i = (\ker A_i)^{\perp\perp} = \text{image } A_{i+1} \oplus H_i$ explains the remaining two. Take U_i to be the orthogonal matrix whose column vectors form the orthonormal basis of the spaces $(\ker A_i)^\perp$ and image A_{i+1} induced from the singular value decomposition of $(\ker A_i)^\perp \to \text{image } A_i$ and $(\ker A_{i+1})^\perp \to \text{image } A_{i+1}$ extended by an orthogonal basis of H_i in the decomposition

$$(\ker A_i)^\perp \oplus \text{image } A_{i+1} \oplus H_i = \mathbb{R}^{c_i}.$$

The linear map A_i has, in terms of these bases, the description $U_{i-1}^t \circ A_i \circ U_i$ which has the desired shape.

Finally, to achieve $\det U_i = 1$, we may, for $1 \leq k \leq r_i$, change signs of the k^{th} column in U_i and $(r_{i-1} + k)^{\text{th}}$ column of U_{i-1} without changing the result of the conjugation. If $h_i > 0$, then changing the sign of any of the last h_i columns of U_i does not affect the result either. Thus, this gives us enough freedom to reach $\det U_i = 1$ for all $i = 0, \ldots, n$. $\qquad\square$

Corollary 2.1 (Repetition of eigenvalues). *Suppose that $A_1, \ldots A_n$ define a complex with $A_0 = A_{n+1} = 0$. Let $\Delta_i = A_i^t \circ A_i + A_{i+1} \circ A_{i+1}^t$ be the corresponding Laplacians. Then, using the orthonormal bases described by the U_i's from Theorem 1.1, the Laplacians have the shape*

$$
\Delta_i =
\begin{array}{c}
 \\
r_i \\
r_{i+1} \\
h_i
\end{array}
\begin{array}{ccc}
r_i & r_{i+1} & h_i \\
\left(\begin{array}{ccc}
\Sigma_i^2 & 0 & 0 \\
0 & \Sigma_{i+1}^2 & 0 \\
0 & 0 & 0
\end{array}\right).
\end{array}
$$

4

In particular,

1. *if $r_i = \operatorname{rank} A_i$ and $\sigma_1^i \geq \sigma_2^i \geq \ldots \sigma_{r_i}^i > 0$ are the singular values of A_i, then each $(\sigma_k^i)^2$ is both an eigenvalue Δ_i and Δ_{i-1};*

2. $\ker \Delta_i = H_i$.

Proof. The structure of Δ_i follows immediately from the structure described in Theorem 1.1. The remaining assertions are immediate consequences. $\square$

Let A_i^+ denote the Moore-Penrose pseudoinverse of the A_i. Thus, a singular value decomposition

$$A_i = U_{i-1} \circ \begin{pmatrix} 0 & 0 & 0 \\ \Sigma_i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \circ U_i^t \quad \text{yields} \quad A_i^+ = U_i \circ \begin{pmatrix} 0 & \Sigma_i^{-1} & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \circ U_{i-1}^t.$$

Proposition 2.2. *Suppose that $A_1, \ldots, A_n$ define a complex with $A_0 = A_{n+1} = 0$. Then, $A_{i+1}^+ \circ A_i^+ = 0$ and*

$$id_{\mathbb{R}^{c_i}} - (A_i^+ \circ A_i + A_{i+1} \circ A_{i+1}^+)$$

defines the orthogonal projection of $\mathbb{R}^{c_i}$ onto the homology H_i.

Proof. We know that $A_i^+ \circ A_i$ defines the projection onto $(\ker A_i)^\perp$ and $A_{i+1} \circ A_{i+1}^+$ defines the projection onto $\operatorname{image} A_{i+1}$. The result follows immediately since these spaces are orthogonal and $H_i = ((\ker A_i)^\perp \oplus \operatorname{image} A_{i+1})^\perp$. $\square$

We call

$$0 \longrightarrow \mathbb{R}^{c_0} \xrightarrow{A_1^+} \mathbb{R}^{c_1} \xrightarrow{A_2^+} \ldots \xrightarrow{A_n^+} \mathbb{R}^{c_n} \longrightarrow 0.$$

the *pseudoinverse complex* of

$$0 \longleftarrow \mathbb{R}^{c_0} \xleftarrow{A_1} \mathbb{R}^{c_1} \xleftarrow{A_2} \ldots \xleftarrow{A_n} \mathbb{R}^{c_n} \longleftarrow 0.$$

Remark 2.3. If the matrices A_i have entries in a subfield $K \subset \mathbb{R}$, then the pseudoinverse complex is also defined over K. This follows since the pseudoinverse is uniquely determined by the Penrose relations [Pen55]:

$$A_i \circ A_i^+ \circ A_i = A_i, \qquad\qquad A_i \circ A_i^+ = (A_i \circ A_i^+)^t,$$

$$A_i^+ \circ A_i \circ A_i^+ = A_i^+, \qquad\qquad A_i^+ \circ A_i = (A_i^+ \circ A_i)^t,$$

which form an algebraic system of equations for the entries of A_i^+ with a unique solution whose coefficients are in K. In particular, this holds for $K = \mathbb{Q}$.

If the entries of the matrices are in the finite field $\mathbb{F}_q$, the pseudoinverse of A_i is well defined over $\mathbb{F}_q$ with respect to the dot-product on $\mathbb{F}_q^{c_i}$ and $\mathbb{F}_q^{c_{i-1}}$ if

$$\ker A_i \cap (\ker A_i)^\perp = 0 \subset \mathbb{F}_q^{c_i} \text{ and } \operatorname{image} A_i \cap (\operatorname{image} A_i)^\perp = 0 \subset \mathbb{F}_q^{c_{i-1}}.$$

We have implemented the computation of the pseudoinverse complex for double precision floating-point numbers $\mathbb{R}_{53}$, the rationals $\mathbb{Q}$, and finite fields $\mathbb{F}_q$ in our Macaulay2 package SVDComplexes.

3   Algorithms

We present two algorithms for computing a singular value decomposition of a complex followed by an example.

Algorithm 3.1 (Successive projection method).
INPUT: A sequences $B_1, \ldots, B_n$ of floating point matrices which are approximations of a complex $A_1, \ldots, A_n$; a threshold b for which we took $b = 10^{-4}$ as default value in our implementation.
OUTPUT: Integers $r_1, \ldots r_n$ and floating point approximations $U_0, \ldots, U_n$ of orthogonal matrices which approximate the singular value decomposition of the corresponding complex.

1. Set $r_0 = 0$, $Q_0 = 0$ and $P_0 = \operatorname{id}_{C_0}$.

2. For i from 1 to n do

    a. Compute the $(c_{i-1} - r_{i-1}) \times c_i$ matrix $\widetilde{B}_i = P_{i-1} \circ B_i$.

    b. Compute the singular value decomposition of $\widetilde{B}_i$, i.e. the diagonal matrix $\widetilde{\Sigma}_i$ of the singular values $\sigma_1^i \geq \sigma_2^i \geq \ldots$ and orthogonal matrices $\widetilde{U}_{i-1}, \widetilde{V}_i^t$ such that
        $$\widetilde{B}_i = \widetilde{U}_{i-1} \circ \widetilde{\Sigma}_i \circ \widetilde{V}_i^t.$$

c. Decide how many singular values of $\widetilde{\Sigma}_i$ are truly non-zero, e.g. for $n_i = \min\{c_{i-1} - r_{i-1}, c_i\}$ and for a threshold b, say $b = 10^{-6}$, take

$$r_i = \begin{cases} \min\{j < n_i \mid b\sigma_j^i \geq \sigma_{j+1}^i\}, & \text{if this set is non-empty} \\ n_i & \text{else} \end{cases}$$

d. Decompose

$$\widetilde{V}_i^t = \begin{pmatrix} Q_i \\ P_i \end{pmatrix}$$

into submatrices consisting of the first r_i and last $c_i - r_i$ rows of $\widetilde{V}_i^t$. (So P_i defines an approximation of the orthogonal projection $C_i \to \ker A_i$ if our guess for r_i was correct.)

e. Define

$$U_{i-1}^t = \begin{pmatrix} Q_{i-1} \\ \widetilde{U}_{i-1}^t \circ P_{i-1} \end{pmatrix}.$$

f. If $i = n$ then $U_n = \widetilde{V}_n^t$.

3. Return $U_0, \ldots, U_n$ and $r_1, \ldots, r_n$.

Proof of concept. We will show that the algorithm gives a good approximation, provided that

  i) the approximation B_i of A_i is sufficiently good,

 ii) we make the correct decisions in step 2.c and

iii) we compute with high enough precision.

By induction on i we will see that P_i defines an approximation of the orthogonal projection $C_i \to \ker A_i$. Since V_i^t is approximately orthogonal

$$\begin{pmatrix} Q_i \\ P_i \end{pmatrix} \circ \begin{pmatrix} Q_i^t & P_i^t \end{pmatrix} \approx \begin{pmatrix} \mathrm{id}_{r_i} & 0 \\ 0 & \mathrm{id}_{c_i - r_i} \end{pmatrix}$$

where id_k denotes a $k \times k$ identity matrix, we additionally conclude that Q_i is an approximation of the orthogonal projection $C_i \to (\ker A_i)^\perp$. This is trivially true in case $i = 0$, since $A_0 = 0$.

For the induction step from $i - 1$ to i, we note that $B_i \approx A_i$ and image $A_i \subset$ ker A_{i-1} implies that $Q_{i-1} \circ B_i \approx 0$. So A_i and $P_{i-1} \circ B_i = \widetilde{B}_i$ have the same 'large' singular values. From

$$\widetilde{B}_i = \widetilde{U}_i \circ \widetilde{\Sigma}_i \circ V_i^t$$

and

$$V_i^t = \begin{pmatrix} Q_i \\ P_i \end{pmatrix}$$

we conclude the assertion that P_i defines approximately the orthogonal projection $C_i \to \ker A_i$ under the assumption, that our choice of r_i is correct. Moreover,

$$U_{i-1}^t \circ A_i \circ U_i \approx U_{i-1}^t \circ B_i \circ U_i$$

$$= \begin{pmatrix} Q_{i-1} \\ \widetilde{U}_{i-1}^t \circ P_{i-1} \end{pmatrix} \circ B_i \circ \begin{pmatrix} Q_i^t & P_i^t \circ \widetilde{U}_i \end{pmatrix}$$

$$\approx \begin{pmatrix} 0 \\ \widetilde{U}_{i-1}^t \circ \widetilde{B}_i \end{pmatrix} \circ \begin{pmatrix} Q_i^t & P_i^t \circ \widetilde{U}_i \end{pmatrix}$$

$$\approx \begin{pmatrix} 0 \\ \widetilde{U}_{i-1}^t \circ \widetilde{U}_{i-1} \circ \widetilde{\Sigma}_i \circ \begin{pmatrix} Q_i \\ P_i \end{pmatrix} \end{pmatrix} \circ \begin{pmatrix} Q_i^t & P_i^t \circ \widetilde{U}_i \end{pmatrix}$$

$$\approx \begin{pmatrix} 0 \\ \widetilde{\Sigma}_i \circ \begin{pmatrix} \mathrm{id}_{r_i} & 0 \\ 0 & \mathrm{id}_{c_i - r_i} \widetilde{U}_i \end{pmatrix} \end{pmatrix}$$

$$\approx \begin{pmatrix} 0 & 0 \\ \Sigma_i & 0 \\ 0 & 0 \end{pmatrix}$$

since

$$\widetilde{\Sigma}_i \circ \begin{pmatrix} 0 \\ \mathrm{id}_{c_i - r_i} \end{pmatrix} \approx 0.$$

This shows that the desired approximation holds. $\qquad\square$

Remark 3.2. To get more confidence in the correctness of the computation of $r_1, \ldots, r_n$ we can alter step 2.c. A natural approach is to start with two approximations $B_1, \ldots, B_n$ and $B'_1, \ldots, B'_n$ in different precisions, and to determine $r_1, \ldots r_n$ as the number of stable singular values, i.e. the singular values which have approximately the same value in both computations.

Our second method quite frequently does not need approximation in two different precisions. It is based on computing with the *Laplacians*

$$\Delta_i = A_i^t \circ A_i + A_{i+1} \circ A_{i+1}^t.$$

Note that $\ker \Delta_i \cong H_i$.

Algorithm 3.3 (Laplacian method).
INPUT: A sequence $B_1, \ldots, B_n$ of floating points approximations of a complex $A_1, \ldots, A_n$, whose Laplacians have no multiple eigenvalues; a threshold b for the relative precision for equality of eigenvalues. In our implementation we took $b = 10^{-4}$ as the default value.
OUTPUT: Integers $r_1, \ldots r_n$ and floating point approximations $U_0, \ldots, U_n$ of orthogonal matrices which approximate the singular value decomposition of the corresponding complex.

1. Compute diagonalisations D_i of the symmetric semi-positive matrices

$$\Delta_i' = B_i^t \circ B_i + B_{i+1} \circ B_{i+1}^t$$

   and orthogonal matrices $U_i' \in SO(c_i)$ such that

$$\Delta_i' = U_i' \circ D_i \circ U_i'^t$$

2. If some D_i has a non-zero eigenvalue with higher multiplicity abort.

3. Let r_i be the number of eigenvalues values which occur up to a chosen relative precision both in D_{i-1} and D_i.

4. Compute the corresponding $c_i \times c_i$ permutation matrices P_i, which put the $r_i + r_{i+1}$ common diagonal entries of D_i into the first positions and set $U_i'' = U_i' \circ P_i$.

5. Compute

$$U_{i-1}''^t \circ B_i \circ U_i'' \approx \begin{pmatrix} 0 & 0 & 0 \\ \Sigma_i' & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

   where the $r_i \times r_i$ matrix Σ_i' is approximately a diagonal matrix.

9

6. Let Σ_i be the diagonal matrix whose entries are the absolute values of the diagonal entries of Σ_i'.

7. Inductively, for i from 1 to n change the signs of the eigenvectors given by the columns of U_i'' to obtain orthogonal matrices U_i such that

$$U_{i-1}^t \circ B_i \circ U_i \approx \begin{pmatrix} 0 & 0 & 0 \\ \Sigma_i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} =: \overline{\Sigma}_i.$$

8. Return $U_0, \ldots, U_n$ and $r_1, \ldots, r_n$.

Proof of concept. We show that we the algorithm produces a good approximation in case

i) the Laplacians $\Delta_i = A_i^t \circ A_i + A_{i+1} \circ A_{i+1}^t$ have no non-zero multiple eigenvalues,

ii) the approximations B_i of the A_i are good enough; in particular, the disturbed non-zero eigenvalues stay apart,

iii) the disturbed zero eigenvalues do not accidentally coincide up to a large relative precision, and

iv) we compute with high enough precision.

Indeed, by ii) and iii) we determine the ranks in step 3 correctly. So in step 5 we will reach approximately a diagonal matrix, and it remains to adjust the signs. $\square$

Remark 3.4. For the choice of the thresholds b in Algorithm 3.1 and Algorithm 3.3 we have only experimental evidence. In particular our default value 10^{-4} has no justification, not even heuristically. We leave it as an open problem to derive a justified choice, which might depend also on the ranks $c_i = \text{rank}\, C_i$ of the $\mathbb{R}$-vector spaces in the complex.

Example 3.5. We consider the complex

$$0 \longleftarrow \mathbb{R}^3 \xleftarrow{A_1} \mathbb{R}^5 \xleftarrow{A_2} \mathbb{R}^5 \xleftarrow{A_3} \mathbb{R}^3 \longleftarrow 0$$

10

defined by the matrices

$$
\begin{pmatrix} 14 & -4 & 16 & 3 & -9 \\ 14 & -5 & 20 & 9 & 1 \\ 4 & 1 & -4 & -12 & -24 \end{pmatrix}, \begin{pmatrix} -43 & -50 & -27 & -51 & 9 \\ 12 & -24 & 36 & 0 & -12 \\ 35 & 34 & 27 & 39 & -9 \\ -3 & -10 & 3 & -6 & -1 \\ -11 & -10 & -9 & -12 & 3 \end{pmatrix}, \begin{pmatrix} -8 & -16 & -12 \\ -5 & -1 & -15 \\ -1 & 13 & -14 \\ 12 & 12 & 28 \\ -1 & 25 & -24 \end{pmatrix}.
$$

The SVD normal form of the complex is given by the matrices

$$
\begin{pmatrix} 34.489 & 0 & 0 & 0 & 0 \\ 0 & 28.714 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 114.08 & 0 & 0 & 0 & 0 \\ 0 & 47.193 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 45.993 & 0 & 0 \\ 0 & 35.209 & 0 \\ 0 & 0 & 0 \end{pmatrix}.
$$

So each of the matrices A_i has rank 2, and all homology groups H_i are 1-dimensional.

The transformation into the normal form is given by the orthogonal matrices

$$
\begin{pmatrix} -.6553 & .2393 & -.7165 \\ -.7549 & -.1745 & .6322 \\ .0262 & .9551 & .2950 \end{pmatrix}, \begin{pmatrix} -.5694 & .1646 & -.7702 & -.1318 & .1950 \\ .1862 & .0303 & .0679 & -.9710 & .1301 \\ -.7448 & -.1213 & .6010 & -.0706 & .2537 \\ -.2631 & -.4289 & -.0790 & -.1821 & -.8411 \\ .1309 & -.8794 & -.1862 & .0404 & .4162 \end{pmatrix},
$$

$$
\begin{pmatrix} .5019 & -.1770 & .2288 & .5338 & .6160 \\ .5257 & .6126 & .3335 & .1127 & -.4738 \\ .3586 & -.7250 & .3461 & -.3015 & -.3677 \\ .5735 & .0970 & -.5972 & -.5061 & .2210 \\ -.1195 & .2417 & .6000 & -.5961 & .4604 \end{pmatrix}, \begin{pmatrix} -.2525 & -.2843 & -.9249 \\ .1813 & -.9528 & .2434 \\ -.9505 & -.1062 & .2921 \end{pmatrix}.
$$

which we have printed here with 4 valid digits only. In other words, the diagram

$$
\begin{array}{ccccccccccc}
0 & \longleftarrow & \mathbb{R}^3 & \xleftarrow{A_1} & \mathbb{R}^5 & \xleftarrow{A_2} & \mathbb{R}^5 & \xleftarrow{A_3} & \mathbb{R}^3 & \longleftarrow & 0 \\
& & \big\uparrow U_0 & & \big\uparrow U_1 & & \big\uparrow U_2 & & \big\uparrow U_3 & & \\
0 & \longleftarrow & \mathbb{R}^3 & \xleftarrow{\overline{\Sigma}_1} & \mathbb{R}^5 & \xleftarrow{\overline{\Sigma}_2} & \mathbb{R}^5 & \xleftarrow{\overline{\Sigma}_3} & \mathbb{R}^3 & \longleftarrow & 0
\end{array}
$$

commutes (up to the chosen precision). The first matrix A_1^+ of the pseudoinverse

complex over $\mathbb{R}_{53}$ printed with 6 valid digits is

$$\begin{pmatrix} .0121907 & .0114627 & .0050431 \\ -.00328525 & -.00426002 & .00115014 \\ .013141 & .0170401 & -.00460058 \\ .00142545 & .00836608 & -.0144655 \\ -.00981498 & .00248076 & -.0291523 \end{pmatrix}$$

which is an approximation of

$$\begin{pmatrix} 5978/490373 & 5621/490373 & 2473/490373 \\ -1611/490373 & -2089/490373 & 564/490373 \\ 6444/490373 & 8356/490373 & -2256/490373 \\ 699/490373 & 8205/980746 & -14187/980746 \\ -4813/490373 & 2433/980746 & -28591/980746 \end{pmatrix}$$

Remark 3.6. This simple example is pretty stable against errors. If we disturb the entries of the matrices in the complex arbitrarily by an relative error of $\leq 10^{-3}$, then taking 10^{-2} as a threshold the algorithms predicts the dimension of the homology groups still correctly, see SVDComplexes.

4   Projection

One application of using the singular value decomposition of a complex is to compute the pseudoinverse complex as described in Section 2. In this section, we consider projecting a sequence of matrices onto a complex.

Algorithm 4.1 (Projection to a complex).
INPUT: A sequence $B_1, \ldots, B_n$ of $c_{i-1} \times c_i$ matrices and a sequence $h_0, \ldots, h_n$ of desired dimension of homology groups.
OUTPUT: A sequence $A_1, \ldots, A_n$ of matrices which define a complex with desired homology, if possible.

1. Set $r_0 = 0$ and compute $r_1, \ldots, r_{n+1}$ from $c_i = r_i + r_{i+1} + h_i$ recursively. If $r_i < 0$ for some i or $r_{n+1} \neq 0$, then return the error message: "The rank conditions cannot be satisfied."

2. Set $Q_0 = 0$ and $P_0 = \mathrm{id}_{C_0}$.

3. For $i = 1, \ldots, n$

    a. Compute the $(c_{i-1} - r_{i-1}) \times c_i$ matrix $\widetilde{B}_i = P_{i-1} \circ B_i$.

    b. Compute the singular value decomposition

$$\widetilde{B}_i = \widetilde{U}_{i-1} \circ \widetilde{\Sigma}_i \circ \widetilde{V}_i^t.$$

    c. Define

$$\overline{\Sigma}_i = \begin{array}{c} \\ r_{i-1} \\ r_i \\ h_{i-1} \end{array} \overset{\begin{array}{ccc} r_i & r_{i+1} & h_i \end{array}}{\left(\begin{array}{ccc} 0 & 0 & 0 \\ \Sigma_i & 0 & 0 \\ 0 & 0 & 0 \end{array} \right)}$$

as a block matrix where Σ_i is a diagonal matrix with entries the largest r_i singular values of $\widetilde{B}_i$.

    d. Decompose

$$\widetilde{V}_i^t = \begin{pmatrix} Q_i \\ P_i \end{pmatrix}$$

into submatrices consisting of the first r_i and last $c_i - r_i$ rows of $\widetilde{V}_i^t$.

    e. Define

$$U_{i-1}^t = \begin{pmatrix} Q_{i-1} \\ \widetilde{U}_{i-1}^t \circ P_{i-1} \end{pmatrix}.$$

    f. If $i = n$, then $U_n = \widetilde{V}_n^t$.

4. Set $A_i = U_{i-1} \circ \overline{\Sigma}_i \circ U_i^t$ and return $A_1, \ldots, A_n$.

Remark 4.2. By construction, it is clear that Algorithm 4.1 computes a complex. We leave it as an open problem to compute the "closest" complex to the given matrices $B_1, \ldots, B_n$.

Example 4.3. In our package RandomComplexes, we have implemented several methods to produce complexes over the integers. The first function randomChain-Complex takes as input a sequences h and r of desired dimension of homology groups and ranks of the matrices. It uses the LLL algorithm [LLL82] to produce example of desired moderate height. It runs fast for complexes of ranks $c_i \leq 100$ but is slow for larger examples because of the use of the LLL-algorithm. Example 3.5 was produced this way.

    We test Algorithms 3.1 and 3.3 to verify the desired dimension of the homology groups. Table 1 compares the timings of these two algorithms on various examples of this sort.

$c_0, \ldots c_3$	$h_0, \ldots, h_3$	Alg. 3.1 (sec)	Alg. 3.3 (sec)
$7, 21, 28, 14$	$2, 3, 2, 1$	.00211	.0110
$8, 27, 35, 17$	$3, 6, 4, 2$	.00225	.0182
$9, 33, 42, 20$	$4, 9, 6, 3$	.00254	.0294
$10, 39, 49, 23$	$5, 12, 8, 4$	.00291	.0647
$11, 45, 56, 26$	$6, 15, 10, 5$	.00355	.1090
$12, 51, 63, 29$	$7, 18, 12, 6$	.00442	.1150

Table 1: Comparison of timings using Algorithms 3.1 and 3.3.

Example 4.4. Our second series of examples is constructed from Stanley-Reisner simplicial complexes of randomly chosen square free monomial ideals. In the specific cases below, we selected N square free monomials at random in a polynomial ring with k variables which are summarized in Table 2. Algorithm 3.3 does not apply to these examples since repeated eigenvalues occur.

5   Application to syzygies

We conclude with an application concerning the computation of Betti numbers in free resolutions. Let $S = K[x_0, \ldots, x_n]$ be the standard graded polynomial ring and M a finitely generated graded S-module. Then, by Hilbert's syzygy theorem, M has a finite free resolution:

$$0 \longleftarrow M \longleftarrow F_0 \xleftarrow{\varphi_1} F_1 \xleftarrow{\varphi_2} \ldots \xleftarrow{\varphi_c} F_c \longleftarrow 0$$

by free graded S-modules $F_i = \sum_j S(-i-j)^{b_{ij}}$ of length $c \le n+1$. Here $S(-\ell)$ denotes the free S-module with generator in degree ℓ.

If we choose in each step a minimal number of homogenous generators, i.e., if $\varphi_i(F_i) \subset (x_0, \ldots x_n) F_{i-1}$, then the free resolution is unique up to an isomorphism. In particular, the Betti numbers b_{ij} of a minimal resolution are numerical invariants of M. On the other hand, for basic applications of free resolutions such as the computation of Ext and Tor-groups, any resolution can be used.

Starting with a reduced Gröbner basis of the submodule $\varphi_1(F_1) \subset F_0$ there is, after some standard choices on orderings, a free resolution such that at each step the columns of φ_{i+1} form a reduced Gröbner basis of $\ker \varphi_i$. This resolution is uniquely determined however, in most cases, highly nonminimal. An algorithm to compute this standard nonminimal resolution was developed in [EMSS16] which

k	N	c_0 h_0	c_1 h_1	c_2 h_2	$\ldots$							Alg 3.1 (sec)
8	20	8	27	44	30							
		1	0	0	1							.00185
9	21	9	35	74	85	46						
		1	0	0	0	0						.0036
10	23	10	45	118	190	173	69					
		1	0	0	0	3	0					.0198
11	26	11	55	165	326	431	361	156	19			
		1	0	0	0	0	0	2	0			.241
12	30	12	66	218	474	694	664	375	101			
		1	0	0	0	0	0	2	0			1.29
13	35	13	78	286	712	1253	1553	1291	639	141		
		1	0	0	0	0	0	0	6	1		39.7
14	41	14	91	364	996	1948	2741	2687	677	559	75	
		1	0	0	0	0	0	0	7	0	0	355.

Table 2: Comparison of timings using Algorithm 3.1.

turned out to be much faster then the computation of a minimal resolution by previous methods.

The following forms the examples which we use as test cases.

Proposition 5.1 (Graded Artinian Gorenstein Algebras). *Let $f \in \mathbb{Q}[x_0, \ldots, x_n]$ be a homogeneous polynomial of degree d. In $S = \mathbb{Q}[\partial_0, \ldots, \partial_n]$, consider the ideal $I = \langle D \in S \mid D(f) = 0 \rangle$ of constant differential operators which annihilate f. Then, $A_f^{\perp} := S/I$ is an artinian Gorenstein Algebra with socle in degree d.*

For more information on this topic see, e.g., [RS00].

Example 5.2. Let $f = \ell_1^4 + \ldots + \ell_{18}^4 \in \mathbb{Q}[x_0, \ldots, x_7]$ be the sum of 4^{th} powers of 18 sufficiently general chosen linear forms ℓ_s. The Betti numbers b_{ij} of the minimal resolution $M = A_f^{\perp}$ as an S-module are zero outside the range $i = 0, \ldots, 8, j = 0, \ldots, 4$. In this range, they take the values:

$j \setminus i$	0	1	2	3	4	5	6	7	8
0	1	.	.	.	.	.	.	.	.
1	.	18	42	.	.	.	.	.	.
2	.	10	63	288	420	288	63	10	.
3	.	.	.	.	.	.	42	18	.
4	.	.	.	.	.	.	.	.	1

which, for example, says that $F_2 = S(-3)^{42} \oplus S(-4)^{63}$. We note that the symmetry of the table is a well-known consequence of the Gorenstein property.

On the other hand the Betti numbers of the uniquely determined nonminimal resolution are much larger:

$j \setminus i$	0	1	2	3	4	5	6	7	8
0	1	.	.	.	.	.	.	.	.
1	.	18	55	75	54	20	3	.	.
2	.	23	145	390	580	515	273	80	10
3	.	7	49	147	245	245	147	49	7
4	.	1	7	21	35	35	21	7	1

To deduce from this resolution the Betti numbers of the minimal resolution, we can use the formula

$$b_{ij} = \dim \operatorname{Tor}_i^S(M, \mathbb{Q})_{i+j}.$$

For example, to deduce $b_{3,2} = 288$, we have to show that the 5^{th} constant strand of the nonminimal resolution

$$0 \longleftarrow \mathbb{Q}^1 \longleftarrow \mathbb{Q}^{49} \longleftarrow \mathbb{Q}^{390} \longleftarrow \mathbb{Q}^{54} \longleftarrow 0$$

has homology only in one position.

The matrices defining the differential in the nonminimal resolution have polynomial entries whose coefficients in $\mathbb{Q}$ can have very large height such that the computation of the homology of the strands becomes infeasible. There are two options, how we can get information about the minimal Betti numbers:

- Pick a prime number p which does not divided any numerator of the normalized reduced Gröbner basis and then reduce modulo p yielding a module $M(p)$ with the same Hilbert function as M. Moreover, for all but finitely many primes p, the Betti numbers of M as an $\mathbb{Q}[x_0, \ldots, x_n]$-module and of $M(p)$ as $\mathbb{F}_p[x_0, \ldots, x_n]$-module coincide.

- Pass from a normalized reduced Gröbner basis of $\varphi_1(F_1) \subset F_0$ to a floating-point approximation of the Gröber basis. Since in the algorithm for the computation of the uniquely determined nonminimal resolution [EMSS16], the majority of ground field operations are multiplications, we can hope that this computation is numerically stable and that the singular value decompositions of the linear strands will detect the minimal Betti numbers correctly.

Example 5.3. We experimented with artinian graded Gorenstein algebras constructed from randomly chosen forms $f \in \mathbb{Q}[x_0, \ldots, x_7]$ in 8 variables which were the sum of n 4^{th} powers of linear forms where $11 \leq n \leq 20$. This experiment showed that roughly 95% the Betti table computed via floating-point arithmetic coincided with one computed over a finite field. The reason for this was that the current implementation uses only double precision floating-point computations which caused difficulty in detecting zero singular values correctly. This would be improved following Remark **??**.

We now consider a series of examples related to the famous Green's conjecture on canonical curves which was proved in a landmark paper [Voi05] for generic curves. In $S = \mathbb{Q}[x_0, \ldots, x_a, y_0, \ldots, y_b]$, consider the homogeneous ideal J_e generated by the 2×2 minors of

$$\begin{pmatrix} x_0 & x_1 & \cdots & x_{a-1} \\ x_1 & x_2 & \cdots & x_a \end{pmatrix} \text{ and } \begin{pmatrix} y_0 & y_1 & \cdots & y_{b-1} \\ y_1 & y_2 & \cdots & y_b \end{pmatrix}$$

together with the entries of the $(a-1) \times (b-1)$ matrix

$$\begin{pmatrix} x_0 & x_1 & x_2 \\ x_1 & x_2 & x_3 \\ \vdots & \vdots & \vdots \\ x_{a-2} & x_{a-1} & x_a \end{pmatrix} \begin{pmatrix} 0 & 0 & e_2 \\ 0 & -e_1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} y_0 & y_1 & \cdots & y_{b-2} \\ y_1 & y_2 & \cdots & y_{b-1} \\ y_2 & y_3 & \cdots & y_b \end{pmatrix}$$

for some parameters $e_1, e_2 \in \mathbb{Q}$. Then, by [ES18], J_e is the homogeneous ideal of an arithmetically Gorenstein surface $X_e(a, b) \subset \mathbb{P}^{a+b+1}$ with trivial canonical bundle. Moreover, the generators of J_e form a Gröbner basis. To verify the generic Green's conjecture for curves of odd genus $g = 2a + 1$, it suffices to prove, for some values $e = (e_1, e_2) \in \mathbb{Q}^2$, that $X_e(a, a)$ has a "natural" Betti table, i.e., for each k there is at most one pair (i, j) with $i + j = k$ and $b_{ij}(X_e(a, a)) \neq 0$. For special values of $e = (e_1, e_2)$, e.g., $e = (0, -1)$, it is known that the resolution is not natural, see [ES18].

Example 5.4. For $a = b = 6$, our implementation computes the following Betti numbers for the nonminimal resolution: as

	0	1	2	3	4	5	6	7	8	9	10	11
1	.	.	.	.	.	.	.	.	.	.	.	.
	.	55	320	930	1688	2060	1728	987	368	81	8	.
	.	.	39	280	906	1736	2170	1832	1042	384	83	8
	.	.	.	1	8	28	56	70	56	28	8	1

For $e = (2, -1)$ and $e = (0, -1)$, our implementation correctly computes the following Betti numbers, respectively, of the minimal resolutions:

	0	1	2	3	4	5	6	7	8	9	10	11
1	1	.	.	.	.	.	.	.	.	.	.	.
	.	55	320	891	1408	1155	.	.	.	.	.	.
	.	.	.	.	.	.	1155	1408	891	320	55	.
	.	.	.	.	.	.	.	.	.	.	.	1

	0	1	2	3	4	5	6	7	8	9	10	11
1	1	.	.	.	.	.	.	.	.	.	.	.
	.	55	320	900	1488	1470	720	315	80	9	.	.
	.	.	9	80	315	720	1470	1488	900	320	55	.
	.	.	.	.	.	.	.	.	.	.	.	1

Each of these computations took several minutes. To consider larger examples, more efficient algorithms and/or implementations for computing the singular value decomposition of a complex are needed.

References

[AFW06] D.N. Arnold, R.S. Falk, and R. Winther, *Finite element exterior calculus, homological techniques, and applications*, Acta Numer. **15** (2006), 1–155.

[AFW10] ———, *Finite element exterior calculus: from Hodge theory to numerical stability*, Bull. Amer. Math. Soc. (N.S.) **47** (2010), no. 2, 281–354.

[EMSS16] B. Eröcal, O. Motsak, F.-O. Schreyer, and A. Steenpaß, *Refined algorithms to compute syzygies*, J. Symbolic Comput. **74** (2016), 308–327.

[ES18] D. Eisenbud and F.-O. Schreyer, *Equation and syzygies of K3 carpets and union of scrolls*, preprint, arXiv:1804.08011 (2018).

[LLL82] A.K. Lenstra, H.W. Lenstra Jr., and L. Lovász, *Factoring polynomials with rational coefficients*, Math. Ann. **261** (1982), no. 4, 515–534.

[M2] D.R. Grayson and M.E. Stillman, *Macaulay2, a software system for research in algebraic geometry*. Available at http://www.math.uiuc.edu/Macaulay2/.

[Pen55] R. Penrose, *A generalized inverse for matrices*, Proc. Cambridge Philos. Soc. **51** (1955), 406–413.

[RS00] K. Ranestad and F.-O. Schreyer, *Varieties of sums of powers*, J. Reine Angew. Math. **525** (2000), 147–181.

[Voi05] C. Voisin, *Green's canonical syzygy conjecture for generic curves of odd genus*, Compos. Math. **141** (2005), no. 5, 1163–1190.

Author Addresses:

Danielle A. Brake
Department of Mathematics, University of Wisconsin - Eau Claire, Eau Claire WI 54702
brakeda@uwec.edu

Jonathan D. Hauenstein
Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, Notre Dame IN 46556
hauenstein@nd.edu

Frank-Olaf Schreyer
Mathematik und Informatik, Universität des Saarlandes, Campus E2 4, D-66123 Saarbrücken, Germany
schreyer@math.uni-sb.de

Andrew J. Sommese
Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, Notre Dame IN 46556
sommese@nd.edu

Michael E. Stillman
Department of Mathematics, Cornell University, Ithaca NY 14853
mike@math.cornell.edu