The Extraordinary SVD

Carla D. Martin and Mason A. Porter

Abstract. The singular value decomposition (SVD) is a popular matrix factorization that has been used widely in applications ever since an efficient algorithm for its computation was developed in the 1970s. In recent years, the SVD has become even more prominent due to a surge in applications and increased computational memory and speed. To illustrate the vitality of the SVD in data analysis, we highlight three of its lesser-known yet fascinating applications. The SVD can be used to characterize political positions of congressmen, measure the growth rate of crystals in igneous rock, and examine entanglement in quantum computation. We also discuss higher-dimensional generalizations of the SVD, which have become increasingly crucial with the newfound wealth of multidimensional data, and have launched new research initiatives in both theoretical and applied mathematics. With its bountiful theory and applications, the SVD is truly extraordinary.

1. IN THE BEGINNING, THERE IS THE SVD. Let's start with one of our favorite theorems from linear algebra and what is perhaps the most important theorem in this paper.

Theorem 1. *Any matrix $A \in \mathbb{R}^{m \times n}$ can be factored into a* singular value decomposition *(SVD),*

$$A = USV^T, \tag{1}$$

where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal matrices (i.e., $UU^T = VV^T = I$) and $S \in \mathbb{R}^{m \times n}$ is diagonal with $r = \text{rank}(A)$ leading positive diagonal entries. The p diagonal entries of S are usually denoted by σ_i for $i = 1, \ldots, p$, where $p = \min\{m, n\}$, and σ_i are called the singular values *of A. The singular values are the square roots of the nonzero eigenvalues of both AA^T and A^TA, and they satisfy the property $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p$.*

See [**66**] for a proof.

Equation (1) can also be written as a sum of rank-1 matrices,

$$A = \sum_{i=1}^{r} \sigma_i u_i v_i^T, \tag{2}$$

where σ_i is the ith singular value, and u_i and v_i are the ith columns of U and V.

Equation (2) is useful when we want to estimate A using a matrix of lower rank [**24**].

Theorem 2 (Eckart-Young). *Let the SVD of A be given by (1). If $k < r = \text{rank}(A)$ and $A_k = \sum_{i=1}^{k} \sigma_i u_i v_i^T$, then*

$$\min_{\text{rank}(B)=k} ||A - B||_2 = ||A - A_k||_2 = \sigma_{k+1}. \tag{3}$$

See [**28**] for a proof.

© THE MATHEMATICAL ASSOCIATION OF AMERICA [Monthly 119

The SVD was discovered over 100 years ago independently by Eugenio Beltrami (1835–1899) and Camille Jordan (1838–1921) [65]. James Joseph Sylvester (1814–1897), Erhard Schmidt (1876–1959), and Hermann Weyl (1885–1955) also discovered the SVD using different methods [65]. The development in the 1960s of practical methods for computing the SVD transformed the field of numerical linear algebra. One method of particular note is the Golub and Reinsch algorithm from 1970 [27]. See [14] for an overview of properties of the SVD and methods for its computation. See the documentation for the Linear Algebra Package (LAPACK) [5] for details on current algorithms to calculate the SVD for dense, structured, or sparse matrices.

Since the 1970s, the SVD has been used in an overwhelming number of applications. The SVD is now a standard topic in many first-year applied mathematics graduate courses and occasionally appears in the undergraduate curriculum. Theorem 2 is one of the most important features of the SVD, as it is extremely useful in least-squares approximations and principal component analysis (PCA). During the last decade, the theory, computation, and application of higher-dimensional versions of the SVD (which are based on Theorem 2) have also become extremely popular among applications with multidimensional data. We include a brief description of a higher-dimensional SVD in this article, and invite you to peruse [37] and references therein for additional details.

We will not attempt in this article to summarize the hundreds of applications that use the SVD, and our discussions and reference list should not be viewed as even remotely comprehensive. Our goal is to summarize a few examples of recent lesser-known applications of the SVD that we enjoy in order to give a flavor of the diversity and power of the SVD, but there are a myriad of others. We mention some of these in passing in the next section, and we then focus on examples from congressional politics, crystallization in igneous rocks, and quantum information theory. We also discuss generalizations of the SVD before ending with a brief summary.

2. IT'S RAINING SVDs (HALLELUJAH)! The SVD constitutes one of science's superheroes in the fight against monstrous data, and it arises in seemingly every scientific discipline.

We find the SVD in statistics in the guise of "principal component analysis" (PCA), which entails computing the SVD of a data set after centering the data for each attribute around the mean. Many other methods of multivariate analysis, such as factor and cluster analysis, have also proven to be invaluable [42]. The SVD per se has been used in chemical physics to obtain approximate solutions to the coupled-cluster equations, which provide one of the most popular tools used for electronic structure calculations [35]. Additionally, we apply an SVD when diagonalizing the one-particle reduced density matrix to obtain the natural orbitals (i.e., the singular vectors) and their occupation numbers (i.e., the singular values). The SVD has also been used in numerous image-processing applications, such as in the calculation of Eigenfaces to provide an efficient representation of facial images in face recognition [50, 68, 69]. It is also important for theoretical endeavors, such as path-following methods for computing curves of equilibria in dynamical systems [23]. The SVD has also been applied in genomics [2, 32], textual database searching [11], robotics [8], financial mathematics [26], compressed sensing [74], and more.

Computing the SVD is expensive for large matrices, but there are now algorithms that offer significant speed-up (see, for example, [10, 40]) as well as randomized algorithms to compute the SVD [41]. The SVD is also the basic structure for higher-dimensional factorizations that are SVD-like in nature [37]; this has transformed computational multilinear algebra over the last decade.

3. CONGRESSMEN ON A PLANE.

In this section, we use the SVD to discuss voting similarities among politicians. In this discussion, we summarize work from [57, 58], which utilize the SVD but focus predominantly on other items.

Mark Twain wrote in *Pudd'nhead Wilson's New Calendar* that "It could probably be shown by facts and figures that there is no distinctly American criminal class except Congress" [70]. Aspects of this snarky comment are actually pretty accurate, as much of the detailed work in making United States law is performed by Congressional committees and subcommittees. (This differs markedly from parliamentary democracies such as Great Britain and Canada.)

There are many ways to characterize the political positions of congressmen. An objective approach is to apply data-mining techniques such as the SVD (or other "multidimensional scaling" methods) on matrices determined by the Congressional Roll Call. Such ideas have been used successfully for decades by political scientists such as Keith Poole and Howard Rosenthal [55, 56]. One question to ask, though, is what observations can be made using just the SVD.

In [57, 58], the SVD was employed to investigate the ideologies of members of Congress. Consider each two-year Congress as a separate data set and also treat the Senate and House of Representatives separately. Define an $m \times n$ voting matrix A with one row for each of the m legislators and one column for each of the n bills on which legislators voted. The element A_{ij} has the value $+1$ if legislator i voted "yea" on bill j and -1 if he or she voted "nay." The sign of a matrix element has no bearing *a priori* on conservativism versus liberalism, as the vote in question depends on the specific bill under consideration. If a legislator did not vote because of absence or abstention, the corresponding element is 0. Additionally, a small number of false zero entries result from resignations and midterm replacements.

Taking the SVD of A allows us to identify congressmen who voted the same way on many bills. Suppose the SVD of A is given by (2). The grouping that has the largest mean-square overlap with the actual groups voting for or against each bill is given by the first left singular vector u_1 of the matrix, the next largest by the second left singular vector u_2, and so on. Truncating A by keeping only the first $k \leq r$ nonzero singular values gives the approximate voting matrix

$$A_k = \sum_{i=1}^{k} \sigma_i u_i v_i^T \approx A. \tag{4}$$

This is a "k-mode truncation" (or "k-mode projection") of the matrix A. By Theorem 2, (4) is a good approximation as long as the singular values decay sufficiently rapidly with increasing i.

A congressman's voting record can be characterized by just two coordinates [57, 58], so the two-mode truncation A_2 is an excellent approximation to A. One of the two directions (the "partisan" coordinate) correlates well with party affiliation for members of the two major parties. The other direction (the "bipartisan" coordinate) correlates well with how often a congressman votes with the majority.[1] We show the coordinates along these first two singular vectors for the 107th Senate (2001–2002) in Figure 1(a). As expected, Democrats (on the left) are grouped together and are almost completely separated from Republicans (on the right).[2] The few instances of party misidentifica-

---

[1] For most Congresses, it suffices to use a two-mode truncation. For a few, it is desirable to keep a third singular vector, which can be used to try to encapsulate a North-South divide [55, 57].

[2] Strictly speaking, the partisanship singular vector is determined up to a sign, which is then chosen to yield the usual Left/Right convention.
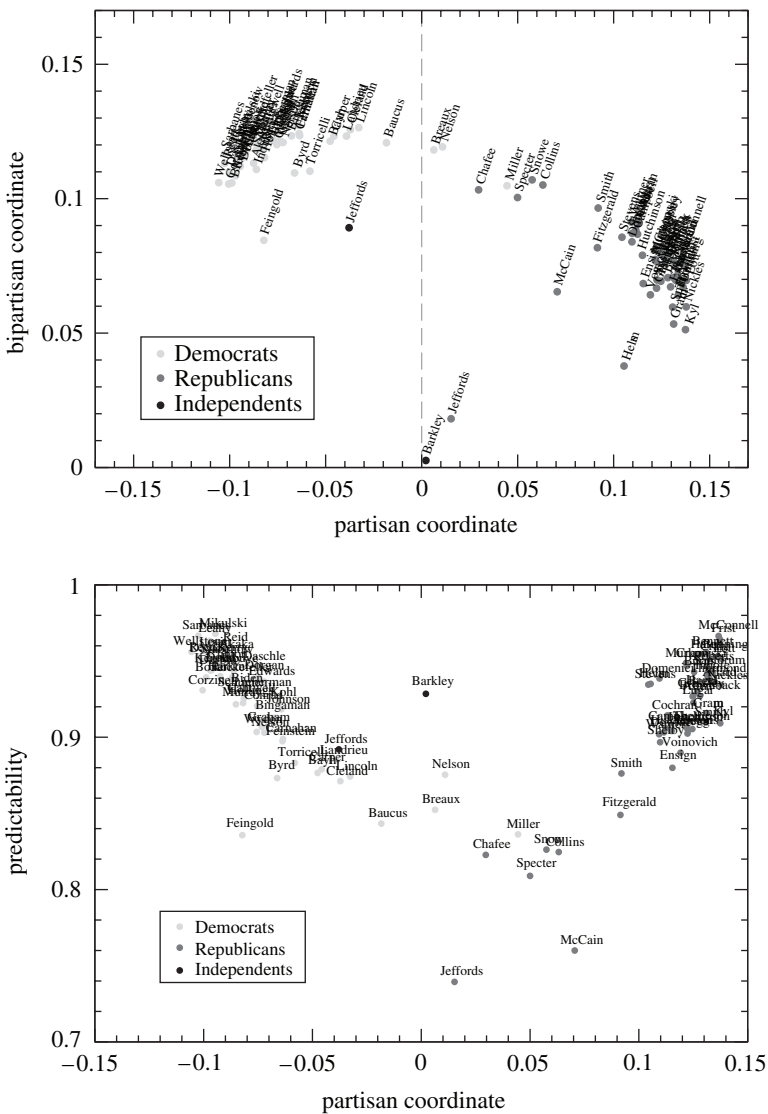
Figure 1. Singular value decomposition (SVD) of the Senate voting record from the 107th U.S. Congress (2001–2002). (a) Two-mode truncation A_2 of the voting matrix A. Each point represents a projection of a single representative's votes onto the leading two eigenvectors (labeled "partisan" and "bipartisan," as explained in the text). Democrats (light dots) appear on the left and Republicans (medium dots) are on the right. The two Independents are shown using dark dots. (b) "Predictability" of votes cast by senators in the 107th Congress based on a two-mode truncation of the SVD. Individual senators range from 74% predictable to 97% predictable. These figures are modified versions of figures that appeared in Ref. [**57**].

tion are unsurprising; conservative Democrats such as Zell Miller [D-GA] appear farther to the right than some moderate Republicans [**12**]. Senator James Jeffords [I-VT], who left the Republican party to become an Independent early in the 107th Congress, appears closer to the Democratic group than the Republican one and to the left of several of the more conservative Democrats.[3]

Equation (4) can also be used to construct an approximation to the votes in the full roll call. Again using A_2, we assign "yea" or "nay" votes to congressmen based on the

---

[3]Jeffords appears twice in Figure 1(a)—once each for votes cast under his two different affiliations.

signs of the matrix elements. Figure 1(b) shows the fraction of actual votes correctly reconstructed using this approximation. Looking at whose votes are easier to reconstruct gives a measure of the "predictability" of the senators in the 107th Congress. Unsurprisingly, moderate senators are less predictable than hard-liners for both parties. Indeed, the two-mode truncation correctly reconstructs the votes of some hard-line senators for as many as 97% of the votes that they cast.

To measure the reproducibility of individual votes and outcomes, the SVD can be used to calculate the positions of the votes along the partisanship and bipartisanship coordinates (see Figure 2). We obtain a score for each vote by reconstituting the voting matrix as before, using the two-mode truncation A_2 and summing the elements of the approximate voting matrix over all legislators. Making a simple assignment of "pass" to those votes that have a positive score and "fail" to all others successfully reconstructs the outcome of 984 of the 990 total votes (about 99.4%) in the 107th House of Representatives. A total of 735 bills passed, so simply guessing that every vote passed would be considerably less effective. This way of counting the success in reconstructing the outcomes of votes is the most optimistic one. Ignoring the values from known absences and abstentions, 975 of the 990 outcomes are still identified correctly. Even the most conservative measure of the reconstruction success rate—in which we ignore values associated with abstentions and absences, assigns individual yeas or nays according to the signs of the elements of A_2, and then observes which outcome has a majority in the resulting roll call—identifies 939 (about 94.8%) of the outcomes correctly. The success rates for other recent Houses are similar [57].
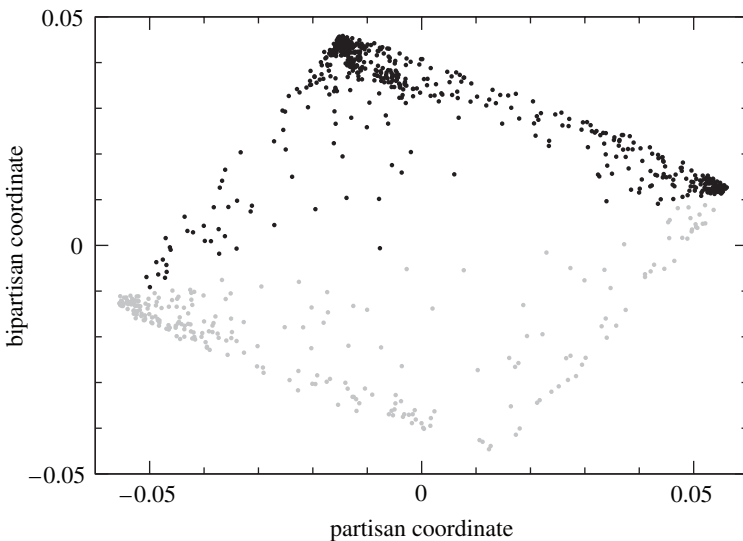

Figure 2. SVD of the roll call of the 107th House of Representatives projected onto the voting coordinates. There is a clear separation between bills that passed (dark dots) and those that did not (light dots). The four corners of the plot are interpreted as follows: bills with broad bipartisan support (north) all passed; those supported mostly by the Right (east) passed because the Republicans were the majority party; bills supported by the Left (west) failed because of the Democratic minority; and the (obviously) very few bills supported by almost nobody (south) also failed. This figure is a modified version of a figure that appeared in [57].

To conclude this section, we remark that it seems to be underappreciated that many political scientists are extremely sophisticated in their use of mathematical and statistical tools. Although the calculations that we discussed above are heuristic ones,

several mathematicians and statisticians have put a lot of effort into using mathematically rigorous methods to study problems in political science. For example, Donald Saari has done a tremendous amount of work on voting methods [20], and (closer to the theme of this article) rigorous arguments from multidimensional scaling have been used recently to study roll-call voting in the House of Representatives [22].

4. THE SVD IS MARVELOUS FOR CRYSTALS.

Igneous rock is formed by the cooling and crystallization of magma. One interesting aspect of the formation of igneous rock is that the microstructure of the rock is composed of interlocking crystals of irregular shapes. The microstructure contains a plethora of quantitative information about the crystallization of deep crust—including the nucleation and growth rate of crystals. In particular, the three-dimensional (3D) *crystal size distribution* (CSD) provides a key piece of information in the study of crystallization rates. CSD can be used, for example, to determine the ratio of nucleation rate to growth rate. Both rates are slow in the deep crust, but the growth rate dominates the nucleation rate. This results in a microstructure composed of large crystals. See [60] for more detail on measuring growth rates of crystals and [31, 43] for more detail on this application of the SVD.

As the crystals in a microstructure become larger, they compete for growth space and their grain shapes become irregular. This makes it difficult to measure grain sizes accurately. CSD analysis of rocks is currently done in two stages. First, take hand measurements of grain sizes in 2D slices and then compute statistical and stereological corrections to the measurements in order to estimate the actual 3D CSD. However, a novel recent approach allows use of the SVD to automatically and directly measure 3D grain sizes that are derived from three specific crystal shapes (prism, plate, and cuboid; see Figure 3) [4]. Ongoing research involves extending such analysis to more complex and irregular shapes. Application to real rock microstructures awaits progress in high energy X-ray tomography, as this will allow improved resolution of grain shapes.
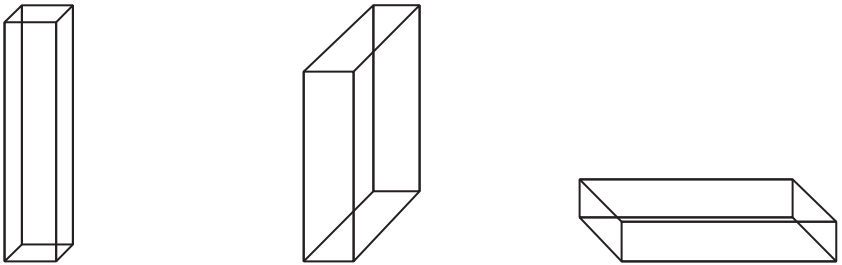
(a) Tetragonal prism (1:1:5)   (b) Tetragonal plate (1:5:5)   (c) Orthorhombic cuboid (1:3:5)

Figure 3. Crystalline structures used to measure grain sizes. We give the relative sizes of their dimensions in parentheses.

The grain sizes are determined by generating databases of microstructures with irregular grain shapes in order to compare the estimated CSD of the actual grains to the computed or ideal CSD predicted by the governing equations. Because the CSDs in many igneous rocks are close to linear [3, 31], the problem can be simplified by using governing equations that generate linear CSDs with the following two rate laws.

1. *Nucleation Rate Law:* $N(t) = e^{\alpha t}$, where N is the number of new nuclei formed at each time step t and α is the nucleation constant.

2. *Crystal Growth Rate Law:* $G = \Delta L/\Delta t$, where $\Delta L/\Delta t$ is the rate of change of a grain diameter per time step. Grain sizes can be represented by short, intermediate, or long diameters. Such diameter classification depends on the relationship between the rate of grain nucleation and the rate of grain growth.

We use an ellipsoid to approximate the size and shape of each grain. There are multiple subjective choices for such ellipsoids that depend on the amount (i.e., the number of points) of the grain to be enclosed by the ellipsoid. To circumvent this subjectivity, it is desirable to compare the results of three types of ellipsoids: the ellipsoid that encloses the entire grain, the ellipsoid that is inscribed within the grain, and the mean of the enclosed and inscribed ellipsoids. See Figure 4 for an illustration of an enclosing and an inscribed ellipsoid.
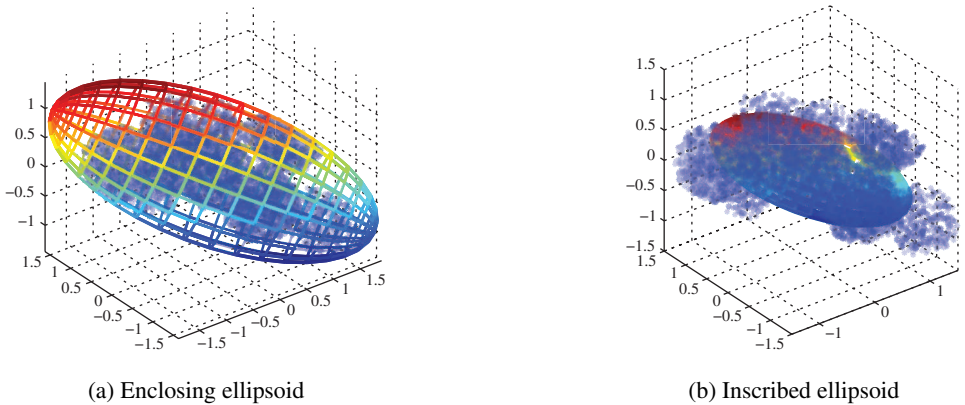

(a) Enclosing ellipsoid   (b) Inscribed ellipsoid

Figure 4. Two possible ellipsoids used to approximate grain sizes. Because grain shapes are irregular, all ellipsoids are triaxial with three unequal diameters.

The SVD is used in the determination of each of the three types of ellipsoids. Comparing the CSDs obtained, using each of the three types of ellipsoids with those predicted by the governing equations, reveals that the inscribed ellipsoids give the best results. In particular, we can use an algorithm developed by Nima Moshtagh [48] that employs the Khachiyan Algorithm [6] along with the SVD to obtain an ellipsoid that encloses an arbitrary number of points (which is defined by the user). Leonid Khachiyan introduced the ellipsoid method in 1979, and this was the first ever worst-case polynomial-time algorithm for linear programming. Given a matrix of data points P containing a discretized set of 3D points representing the crystal, we solve

$$\min_{A,c} \log\{\det(A)\} \quad \text{subject to} \quad (P_i - c)^T A (P_i - c) \le 1, \tag{5}$$

where P_i is the ith column of P, the matrix A contains information about the shape of the ellipsoid, and c is the center of the ellipsoid.

Note that P in this case is dense, it has size $n \times 3$, and $n \approx 5000$. Once A and c have been determined, we calculate the ith radius of the D-dimensional ellipse from the SVD of A using

$$r_i = 1/\sqrt{\sigma_i}, \tag{6}$$

where σ_i $(i = 1, \ldots, D)$ is the ith singular value of A. If the SVD of A is given by equation (1), then the orientation of the ellipsoid is given by the rotation matrix V.

The major difficulty in such studies of igneous rock is that grain shapes and sizes are irregular due to competition for growth space among crystals. In particular, they are not of the ideal sizes and shapes that are assumed by crystallization theory. For example, crystals might start to grow with definite diameter ratios (yielding, for example, the prism, plate, or cuboid in Figure 3) but eventually develop irregular outlines. Current studies [**4**] suggest that one of the diameters or radii of the inscribed ellipsoid (as determined from the SVD) can be used as a measure of grain size for the investigation of crystal size distributions, but the problem remains open.

5. QUANTUM INFORMATION SOCIETY. From a physical perspective, information is encoded in the state of a physical system, and a computation is carried out on a physically realizable device [**59**]. *Quantum information* refers to information that is held in the state of a quantum system. Research in quantum computation and quantum information theory has helped lead to a revival of interest in linear algebra by physicists. In these studies, the SVD (especially in the form of the Schmidt decomposition) have been crucial for gaining a better understanding of fundamental quantum-mechanical notions such as entanglement and measurement.

Entanglement is a quantum form of correlation that is much stronger than classical correlation, and quantum information scientists use entanglement as a basic resource in the design of quantum algorithms [**59**]. The potential power of quantum computation relies predominantly on the inseparability of multipartite quantum states, and the extent of such interlocking can be measured using entanglement.

We include only a brief discussion in the present article, but one can go much farther [**54, 59, 62**]. Whenever there are two distinguishable particles, we can fully characterize inseparable quantum correlations using what is known as a "single-particle reduced density matrix" (see the definition below), and the SVD is crucial for demonstrating that this is the case. See [**54, 59, 62**] for lots of details and all of the quantum mechanics notation that you'll ever desire.

Suppose that we have two distinguishable particles A and B. We can then write a joint pure-state wave function $|\Psi\rangle$, which is expressed as an expansion in its states weighted by the probability that they occur. Note that we have written the wave function using Dirac (bra-ket) notation. It is a column vector, and its Hermitian conjugate is the row vector $\langle\Psi|$. The prefactor for each term in the expansion of $|\Psi\rangle$ consists of the complex-valued components C_{ij} of an $m \times n$ probability matrix C, which satisfies $\mathrm{tr}(CC^{\dagger}) = \mathrm{tr}(C^{\dagger}C) = 1$. (Recall that $X^{\dagger}$ refers to the Hermitian conjugate of the matrix X.)

Applying the SVD of C (i.e., letting $C = USV^{\dagger}$, where U and V are unitary matrices[4]) and transforming to a single-particle basis allows us to diagonalize $|\Psi\rangle$, which is said to be *entangled* if more than one singular value is nonzero. We can even measure the entanglement using the two-particle density matrix $\rho := |\Psi\rangle\langle\Psi|$ that is given by the outer product of the wave function with itself. We can then compute the von Neumann entanglement entropy

$$\sigma = -\sum_{k=1}^{\min(n,m)} |S_k^2| \ln |S_k^2|. \tag{7}$$

Because $|S_k^2| \in [0, 1]$, the entropy is zero for unentangled states and has the value $\ln[\min(n, m)]$ for maximally entangled states.

---

[4]A unitary matrix U satisfies $UU^{\dagger} = 1$ and is the complex-valued generalization of an orthogonal matrix.

The SVD is also important in other aspects of quantum information. For example, it can be used to help construct measurements that are optimized to distinguish between a set of (possibly nonorthogonal) quantum states [25].

6. CAN YOU TAKE ME HIGHER? As we have discussed, the SVD permeates numerous applications and is vital to data analysis. Moreover, with the availability of cheap memory and advances in instrumentation and technology, it is now possible to collect and store enormous quantities of data for science, medical, and engineering applications. A byproduct of this wealth is an ever-increasing abundance of data that is fundamentally three-dimensional or higher. The information is thus stored in multiway arrays—i.e., as tensors—instead of as matrices. An order-p tensor $\mathcal{A}$ is a multiway array with p indices:

$$\mathcal{A} = (a_{i_1 i_2 \dots i_p}) \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_p}.$$

Thus, a first-order tensor is a vector, a second-order tensor is a matrix, a third-order tensor is a "cube," and so on. See Figure 5 for an illustration of a $2 \times 2 \times 2$ tensor.
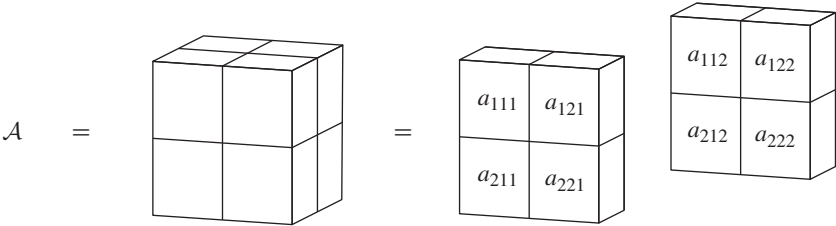

Figure 5. Illustration of a $2 \times 2 \times 2$ tensor as a cube of data. This figure originally appeared in [34] and is used with permission from Elsevier.

Applications involving operations with tensors are now widespread. They include chemometrics [64], psychometrics [38], signal processing [15, 17, 63], computer vision [71, 72, 73], data mining [1, 61], networks [36, 49], neuroscience [7, 46, 47], and many more. For example, the facial recognition algorithm *Eigenfaces* [50, 68, 69] has been extended to *TensorFaces* [71]. To give another example, experiments have shown that fluorescence (i.e., the emission of light from a substance) is modeled well using tensors, as the data follow a trilinear model [64].

A common thread in these applications is the need to manipulate the data, usually by compression, by taking advantage of its multidimensional structure (see, for example, the recent article [52]). Collapsing multiway data to matrices and using standard linear algebra to answer questions about the data often has undesirable consequences. It is thus important to consider the multiway data directly.

Here we provide a brief overview of two types of higher-order extensions of the matrix SVD. For more information, see the extensive article on tensor decompositions [37] and references therein. Recall from (2) that the SVD is a rank-revealing decomposition. The outer product $u_i v_i^T$ in equation (2) is often written using the notation $u_i \circ v_i$. Just as the outer product of two vectors is a rank-1 matrix, the outer product of three vectors is a rank-1 third-order tensor. For example, if $x \in \mathbb{R}^{n_1}$, $y \in \mathbb{R}^{n_2}$, and $z \in \mathbb{R}^{n_3}$, then the outer product $x \circ y \circ z$ has dimension $n_1 \times n_2 \times n_3$ and is a rank-1 third-order tensor whose (i, j, k)th entry is given by $x_i y_j z_k$. Likewise, an outer product of four vectors gives a rank-1 fourth-order tensor, and so on. For the rest of this discussion,

we will limit our exposition to third-order tensors, but the concepts generalize easily to order-p tensors.

The *tensor rank r* of an order-p tensor $\mathcal{A}$ is the minimum number of rank-1 tensors that are needed to express the tensor. For a third-order tensor $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, this implies the representation

$$\mathcal{A} = \sum_{i=1}^{r} \sigma_i (u_i \circ v_i \circ w_i), \tag{8}$$

where σ_i is a scaling constant. The scaling constants are the nonzero elements of an $r \times r \times r$ diagonal tensor $S = (\sigma_{ijk})$. (As discussed in [**37**], a tensor is called *diagonal* if the only nonzero entries occur in elements σ_{ijk} with $i = j = k$.) The vectors u_i, v_i, and w_i are the ith columns from matrices $U \in \mathbb{R}^{n_1 \times r}$, $V \in \mathbb{R}^{n_2 \times r}$, and $W \in \mathbb{R}^{n_3 \times r}$, respectively.

We can think of equation (8) as an extension of the matrix SVD. Note, however, the following differences.

1. The matrices U, V, and W in (8) are *not* constrained to be orthogonal. Furthermore, an orthogonal decomposition of the form (8) does not exist, except in very special cases [**21**].

2. The maximum possible rank of a tensor is not given directly from the dimensions, as is the case with matrices.[5] However, loose upper bounds on rank do exist for higher-order tensors. Specifically, the maximum possible rank of an $n_1 \times n_2 \times n_3$ tensor is bounded by $\min(n_1 n_2, n_1 n_3, n_2 n_3)$ in general [**39**] and $\lfloor 3n/2 \rfloor$ in the case of $n \times n \times 2$ tensors [**9, 33, 39, 44**]. In practice, however, the rank is typically much less than these upper bounds. For example, [**16**] conjectures that the rank of a particular $9 \times 9 \times 9$ tensor is 19 or 20.

3. Recall that the best rank-k approximation to a matrix is given by the kth partial sum in the SVD expansion (Theorem 2). However, this result does not extend to higher-order tensors. In fact, the best rank-k approximation to a tensor might not even exist [**19, 53**].

4. There is no known closed-form solution to determine the rank r of a tensor *a priori*; in fact, the problem is NP-hard [**30**]. Rank determination of a tensor is a widely-studied problem [**37**].

In light of these major differences, there exists more than one higher-order version of the matrix SVD. The different available decompositions are motivated by the application areas. A decomposition of the form (8) is called a CANDECOMP-PARAFAC (CP) decomposition (CANonical DECOMPosition or PARAllel FACtors model) [**13, 29**], whether or not r is known to be minimal. However, since an orthogonal decomposition of the form (8) does not always exist, a *Tucker3* form is often used to guarantee the existence of an orthogonal decomposition as well as to better model certain data [**51, 61, 71, 72, 73**].

If $\mathcal{A}$ is an $n_1 \times n_2 \times n_3$ tensor, then its *Tucker3 decomposition* has the form [**67**]

$$\mathcal{A} = \sum_{i=1}^{m_1} \sum_{j=1}^{m_2} \sum_{k=1}^{m_3} \sigma_{ijk} (u_i \circ v_j \circ w_k), \tag{9}$$

---

[5]The maximum possible rank of an $n_1 \times n_2$ matrix is $\min(n_1, n_2)$.

where u_i, v_j, and w_k are the ith, jth, and kth columns of the matrices $U \in \mathbb{R}^{n_1 \times m_1}$, $V \in \mathbb{R}^{n_2 \times m_2}$, and $W \in \mathbb{R}^{n_3 \times m_3}$. Often, U, V, and W have orthonormal columns. The tensor $S = (\sigma_{ijk}) \in \mathbb{R}^{m_1 \times m_2 \times m_3}$ is called the *core tensor*. In general, the core tensor S is dense and the decomposition (9) does not reveal its rank. Equation (9) has also been called the higher-order SVD (HOSVD) [**18**], though the term "HOSVD" actually refers to a method for computation [**37**]. Reference [**18**] demonstrates that the HOSVD is a convincing extension of the matrix SVD. The HOSVD is guaranteed to exist, and it computes (9) directly by calculating the SVDs of the three matrices obtained by "flattening" the tensor into matrices in each dimension and then using those results to assemble the core tensor. Yet another extension of the matrix SVD factors a tensor as a product of tensors rather than as an outer product of vectors [**34, 45**].

7. EVERYWHERE YOU GO, ALWAYS TAKE THE SVD WITH YOU. The SVD is a fascinating, fundamental object of study that has provided a great deal of insight into a diverse array of problems, which range from social network analysis and quantum information theory to applications in geology. The matrix SVD has also served as the foundation from which to conduct data analysis of multiway data by using its higher-dimensional tensor versions. The abundance of workshops, conference talks, and journal papers in the past decade on multilinear algebra and tensors also demonstrates the explosive growth of applications for tensors and tensor SVDs. The SVD is an omnipresent factorization in a plethora of application areas. We recommend it highly.

REFERENCES

1. E. Acar, S. A. Çamtepe, M. S. Krishnamoorthy, B. Yener, Modeling and multiway analysis of chatroom tensors, in *Intelligence and Security Informatics, Lecture Notes in Computer Science*, Vol. 3495, 181–199. Kantor, Paul and Muresan, Gheorghe and Roberts, Fred and Zeng, Daniel and Wang, Fei-Yue and Chen, Hsinchun and Merkle, Ralph, eds. Springer, Berlin/Heidelberg, 2005.
2. O. Alter, P. O. Brown, D. Botstein, Singular value decomposition for genome-wide expression data processing and modeling, *Proceedings of the National Academy of Sciences* **97** (2000) 10101–10106, available at `http://dx.doi.org/10.1073/pnas.97.18.10101`.
3. R. Amenta, A. Ewing, A. Jensen, S. Roberts, K. Stevens, M. Summa, S. Weaver, P. Wertz, A modeling approach to understanding the role of microstructure development on crystal-size distributions and on recovering crystal-size distributions from thin slices, *American Mineralogist* **92** (2007) 1936–1945, available at `http://dx.doi.org/10.2138/am.2007.2408`.
4. R. Amenta, B. Wihlem, Application of singular value decomposition to estimating grain sizes for crystal size distribution analysis, *GAC-MAC-SEG-SGA Ottawa 2011*, available at `http://www.mineralogicalassociation.ca/index.php?p=35`.
5. E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, *LAPACK User's Guide*, third edition. Society for Industrial and Applied Mathematics, Philadelphia, 1999.
6. B. Aspvall, R. E. Stone, Khachiyan's linear programming algorithm, *Journal of Algorithms* **1** (1980) 1–13, available at `http://dx.doi.org/10.1016/0196-6774(80)90002-4`.
7. C. Beckmann, S. Smith, Tensorial extensions of the independent component analysis for multisubject fMRI analysis, *NeuroImage* **25** (2005) 294–311, available at `http://dx.doi.org/10.1016/j.neuroimage.2004.10.043`.
8. C. Belta, V. Kumar, An SVD-based projection method for interpolation on $SE(3)$, *IEEE Transactions on Robotics and Automation* **18** (2002) 334–345, available at `http://dx.doi.org/10.1109/TRA.2002.1019463`.

9. J. M. F. ten Berge, Kruskal's polynomial for $2 \times 2 \times 2$ arrays and a generalization to $2 \times n \times n$ arrays, *Psychometrika* **56** (1991) 631–636, available at `http://dx.doi.org/10.1007/BF02294495`.

10. M. W. Berry, Large scale sparse singular value computations, *International Journal of Supercomputer Applications* **6** (1992) 13–49.

11. M. W. Berry, S. T. Dumais, G. W. O'Brien, Using linear algebra for intelligent information retrieval, *SIAM Review* **37** (1995) 573–595, available at `http://dx.doi.org/10.1137/1037127`.

12. J. R. Boyce, D. P. Bischak, The role of political parties in the organization of Congress, *The Journal of Law, Economics, & Organization* **18** (2002) 1–38.

13. J. D. Carroll, J. Chang, Analysis of individual differences in multidimensional scaling via an *N*-way generalization of "Eckart-Young" decomposition, *Psychometrika* **35** (1970) 283–319, available at `http://dx.doi.org/10.1007/BF02310791`.

14. A. K. Cline, I. S. Dhillon, Computation of the singular value decomposition, in *Handbook of Linear Algebra*. Edited by L. Hogben. CRC Press, Boca Raton, FL, 2006, 45.1–45.13.

15. P. Comon, Tensor decompositions, in *Mathematics in Signal Processing V*. Edited by J. G. McWhirter and I. K. Proudler. Clarendon Press, Oxford, 2002, 1–24.

16. P. Comon, J. M. F. ten Berge, L. De Lathauwer, J. Castaing, Generic and typical ranks of multi-way arrays, *Linear Algebra and its Applications* **430** (2009) 2997–3007, available at `http://dx.doi.org/10.1016/j.laa.2009.01.014`.

17. L. De Lathauwer, B. De Moor, From matrix to tensor: Multilinear algebra and signal processing, in *Mathematics in Signal Processing IV*. Editd by J. McWhirter and I.K. Proudler. Clarendon Press, Oxford, 1998, 1–15.

18. L. De Lathauwer, B. De Moor, J. Vandewalle, A multilinear singular value decomposition, *SIAM Journal of Matrix Analysis and Applications* **21** (2000) 1253–1278.

19. V. De Silva, L.-H. Lim, Tensor rank and the ill-posedness of the best low-rank approximation problem, *SIAM Journal on Matrix Analysis and Applications* **30** (2008) 1084–1127.

20. *Decisions and Elections: Explaining the Unexpected*. Edited by D. G. Saari. Cambridge University Press, Cambridge, 2001.

21. J. B. Denis, T. Dhorne, Orthogonal tensor decomposition of 3-way tables, in *Multiway Data Analysis*. Edited by R. Coppi and S. Bolasco. Elsevier, Amsterdam, 1989, 31–37.

22. P. Diaconis, S. Goel, S. Holmes, Horseshoes in multidimensional scaling and local kernel methods, *Annals of Applied Statistics* **2** (2008) 777–807, available at `http://dx.doi.org/10.1214/08-AOAS165`.

23. L. Dieci, M. G. Gasparo, A. Papini, Path following by SVD, in *Computational Science - ICCS 2006, Lecture Notes in Computer Science*, Vol. 3994, 2006, 677–684.

24. G. Eckart, G. Young, The approximation of one matrix by another of lower rank, *Psychometrika* **1** (1936) 211–218, available at `http://dx.doi.org/10.1007/BF02288367`.

25. Y. C. Eldar, G. D. Forney, Jr., On quantum detection and the square-root measurement, *IEEE Transactions on Information Theory* **47** (2001) 858–872, available at `http://dx.doi.org/10.1109/18.915636`.

26. D. J. Fenn, M. A. Porter, S. Williams, M. McDonald, N. F. Johnson, N. S. Jones, Temporal evolution of financial market correlations, *Physical Review E* **84** (2011) 026109, available at `http://dx.doi.org/10.1103/PhysRevE.84.026109`.

27. G. H. Golub, C. Reinsch, Singular value decomposition and least squares solutions, *Numerische Mathematik* **14** (1970) 403–420, available at `http://dx.doi.org/10.1007/BF02163027`.

28. G. H. Golub, C. F. Van Loan, *Matrix Computations*, third edition, The Johns Hopkins University Press, Baltimore, MD, 1996.

29. R. A. Harshman, Foundations of the PARAFAC procedure: Model and conditions for an 'explanatory' multi-mode factor analysis, *UCLA Working Papers in Phonetics* **16** (1970) 1–84.

30. J. Hastad, Tensor rank is NP-complete, *Journal of Algorithms* **11** (1990) 654–664, available at `http://dx.doi.org/10.1016/0196-6774(90)90014-6`.

31. M. D. Higgins, *Quantitative Textural Measurements in Igneous and Metamorphic Petrology*, Cambridge University Press, Cambridge, U.K., 2006.

32. N. S. Holter, M. Mitra, A. Maritan, M. Cieplak, J. R. Banavar, N. V. Fedoroff, Fundamental patterns underlying gene expression profiles: Simplicity from complexity, *Proceedings of the National Academy of Sciences* **97** (2000) 8409–8414, available at `http://dx.doi.org/10.1073/pnas.150242097`.

33. J. Ja'Ja', Optimal evaluation of pairs of bilinear forms, *SIAM Journal on Computing* **8** (1979) 443–461, available at `http://dx.doi.org/10.1137/0208037`.

34. M. E. Kilmer, C. D. Martin, Factorization strategies for third-order tensors, *Linear Algebra and its Applications* **435** (2011) 641–658, available at `http://dx.doi.org/10.1016/j.laa.2010.09.020`.

35. T. Kinoshita, Singular value decomposition approach for the approximate coupled-cluster method, *Journal of Chemical Physics* **119** (2003) 7756–7762, available at `http://dx.doi.org/10.1063/1.1609442`.

36. T. G. Kolda, B. W. Bader, Higher-order web link analysis using multilinear algebra, in *Data Mining*

ICDM 2005, Proceedings of the 5th IEEE International Conference, IEEE Computer Society, 2005, 242–249.

37. ———, Tensor decompositions and applications, *SIAM Review* **51** (2009) 455–500, available at `http://dx.doi.org/10.1137/07070111X`.

38. P. M. Kroonenberg, *Three-Mode Principal Component Analysis: Theory and Applications*, DSWO Press, Leiden, 1983.

39. J. B. Kruskal, Rank, decomposition, and uniqueness for 3-way and N-way arrays, in *Multiway Data Analysis*, Edited by R. Coppi and S. Bolasco, Elsevier, Amsterdam, 1989, 7–18.

40. R. M. Larsen, Lanczos bidiagonalization with partial reorthogonalization, in *Efficient Algorithms for Helioseismic Inversion*, Part II, Chapter A, Ph.D. Thesis, Department of Computer Science, University of Aarhus, 1998.

41. E. Liberty, F. Woolfe, P.-G. Martinsson, V. Rokhlin, M. Tygert, Randomized algorithms for the low-rank approximation of matrices, *Proceedings of the National Academy of Sciences* **104** (2007) 20167–20172, available at `http://dx.doi.org/10.1073/pnas.0709640104`.

42. K. Marida, J. T. Kent, S. Holmes, *Multivariate Analysis*, Academic Press, New York, 2005.

43. A. D. Marsh, On the interpretation of crystal size distributions in magmatic systems, *Journal of Petrology* **39** (1998) 553–599, available at `http://dx.doi.org/10.1093/petroj/39.4.553`.

44. C. D. Martin, The rank of a $2 \times 2 \times 2$ tensor, *Linear and Multilinear Algebra*, **59** (2011) 943–950, available at `http://dx.doi.org/10.1080/03081087.2010.538923`.

45. C. D. Martin, R. Shafer, B. LaRue, A recursive idea for multiplying order-p tensors, *SIAM Journal of Scientific Computing*, to appear (2012).

46. E. Martínez-Montes, P. A. Valdés-Sosa, F. Miwakeichi, R. I. Goldman, M. S. Cohen, Concurrent EEG/fMRI analysis by multiway partial least squares, *NeuroImage* **22** (2004) 1023–1034, available at `http://dx.doi.org/10.1016/j.neuroimage.2004.03.038`.

47. F. Miwakeichi, E. Martínez-Montes, P. A. Valdés-Sosa, N. Nishiyama, H. Mizuhara, Y. Yamaguchi, Decomposing EEG data into space-time-frequency components using parallel factor analysis, *NeuroImage* **22** (2004) 1035–1045, available at `http://dx.doi.org/10.1016/j.neuroimage.2004.03.039`.

48. N. Moshtagh, Minimum volume enclosing ellipsoid. From MATLAB Central File Exchange (1996) MathWorks, available at `http://www.mathworks.com/matlabcentral/fileexchange/9542-minimum-volume-enclosing-ellipsoid`.

49. P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, J.-P. Onnela, Community structure in time-dependent, multiscale, and multiplex networks, *Science* **328** (2010) 876–878, available at `http://dx.doi.org/10.1126/science.1184819`.

50. N. Muller, L. Magaia, B. M. Herbst, Singular value decomposition, eigenfaces, and 3D reconstructions, *SIAM Review* **46** (2004) 518–545, available at `http://dx.doi.org/10.1137/S0036144501387517`.

51. J. G. Nagy, M. E. Kilmer, Kronecker product approximation for preconditioning in three-dimensional imaging applications, *IEEE Trans. Image Proc.* **15** (2006) 604–613, available at `http://dx.doi.org/10.1109/TIP.2005.863112`.

52. I. V. Oseledets, D. V. Savostianov, E. E. Tyrtyshnikov, Tucker dimensionality reduction of three-dimensional arrays in linear time, *SIAM Journal of Matrix Analysis and Applications* **30** (2008) 939–956, available at `http://dx.doi.org/10.1137/060655894`.

53. P. Paatero, Construction and analysis of degenerate PARAFAC models, *Journal of Chemometrics*, **14** (2000) 285–299.

54. R. Paškauskas, L. You, Quantum correlations in two-boson wave functions, *Physical Review A* **64** (2001) 042310, available at `http://dx.doi.org/10.1103/PhysRevA.64.042310`.

55. K. T. Poole, Voteview, Department of Political Science, University of Georgia, 2011, available at `http://voteview.com`.

56. K. T. Poole, H. Rosenthal, *Congress: A Political-Economic History of Roll Call Voting*, Oxford University Press, Oxford, 1997.

57. M. A. Porter, P. J. Mucha, M. E. J. Newman, A. J. Friend, Community structure in the United States House of Representatives, *Physica A* **386** (2007) 414–438, available at `http://dx.doi.org/10.1016/j.physa.2007.07.039`.

58. M. A. Porter, P. J. Mucha, M. E. J. Newman, C. M. Warmbrand, A network analysis of committees in the United States House of Representatives, *Proceedings of the National Academy of Sciences* **102** (2005) 7057–7062, available at `http://dx.doi.org/10.1073/pnas.0500191102`.

59. J. Preskill, Lecture notes for physics 229: Quantum information and computation (2004), California Institute of Technology, available at `http://www.theory.caltech.edu/~preskill/ph229/`.

60. A. Randolph, M. Larson, *Theory of Particulate Processes*, Academic Press, New York, 1971.

61. B. Savas, L. Eldén, Handwritten digit classification using higher-order singular value decomposition, *Pattern Recognition* **40** (2007) 993–1003, available at `http://dx.doi.org/10.1016/j.patcog.2006.08.004`.

62. J. Schliemann, D. Loss, A. H. MacDonald, Double-occupancy errors, adiabaticity, and entanglement of spin qubits in quantum dots, *Physical Review B* **63** (2001) 085311, available at `http://dx.doi.org/10.1103/PhysRevB.63.085311`.

63. N. Sidiropoulos, R. Bro, G. Giannakis, Parallel factor analysis in sensor array processing, *IEEE Transactions on Signal Processing* **48** (2000) 2377–2388, available at `http://dx.doi.org/10.1109/78.852018`.

64. A. Smilde, R. Bro, P. Geladi, *Multi-way Analysis: Applications in the Chemical Sciences*, Wiley, Hoboken, NJ, 2004.

65. G. W. Stewart, On the early history of the singular value decomposition, *SIAM Review* **35** (1993) 551–566, available at `http://dx.doi.org/10.1137/1035134`.

66. G. Strang, *Linear Algebra and its Applications*, fourth edition, Thomson Brooks/Cole, 2005.

67. L. R. Tucker, Some mathematical notes on three-mode factor analysis, *Psychometrika* **31** (1966) 279–311, available at `http://dx.doi.org/10.1007/BF02289464`.

68. M. Turk, A. Pentland, Eigenfaces for recognition, *Journal of Cognitive Neuroscience* **3** (1991a), available at `http://dx.doi.org/10.1162/jocn.1991.3.1.71`.

69. ———— , Face recognition using Eigenfaces, *Proc. of Computer Vision and Pattern Recognition* **3** (1991b) 586–591, available at `http://dx.doi.org/10.1109/CVPR.1991.139758`.

70. M. Twain, Pudd'nhead Wilson's new calendar, in *Following the Equator*, Samuel L. Clemens, Hartford, CT, 1897.

71. M. A. O. Vasilescu, D. Terzopoulos, Multilinear analysis of image ensembles: Tensorfaces, *Computer Vision – ECCV 2002*, Proceedings of the 7th European Conference, Lecture Notes in Computer Science, Vol. 2350, 2002, 447–460.

72. ———— , Multilinear image analysis for face recognition, *Pattern Recognition–ICPR 2002, Proceedings of the International Conference*, Vol. 2, 2002, 511–514.

73. ———— , Multilinear subspace analysis of image ensembles, *Computer Vision and Pattern Recognition–CVPR 2003*, Proceedings of the 2003 IEEE Computer Society Conference, 2003, 93–99.

74. L. Xu, Q. Liang, Computation of the singular value decomposition, in *Wireless Algorithms, Systems, and Applications 2010, Lecture Notes in Computer Science*, Vol. 6221, Edited by G. Pandurangan, V. S. A. Kumar, G. Ming, Y. Liu, and Y. Li, Springer-Verlag, Berlin, 2010, 338–342.

CARLA D. MARTIN is an associate professor of mathematics at James Madison University. She has a strong bond with linear algebra and especially with the SVD, inspired by her thesis advisor at Cornell University, Charles Van Loan. She performs research in multilinear algebra and tensors but pretty much loves anything having to do with matrices. She has been a huge promoter of publicizing mathematical applications in industry, both in talks and in print, and also serves as the VP of Programs in BIG SIGMAA. She is an active violinist and enjoys teaching math and music to her three children.

She wishes to especially recognize the late Roddy Amenta, whose work is highlighted in Section 4. Roddy passed away while this article was in press. He was a powerful educator and was very dedicated to his students and to the field of geology.

Department of Mathematics and Statistics, James Madison University, Harrisonburg, VA 22807
carla.dee@gmail.com

MASON A. PORTER is a University Lecturer in the Mathematical Institute at the University of Oxford. He is a member of the Oxford Centre for Industrial and Applied Mathematics (OCIAM) and of the CABDyN Complexity Centre. He is also a Tutorial Fellow of Somerville College in Oxford. Mason's primary research areas are nonlinear and complex systems, but at some level he is interested in just about everything. Mason originally saw the SVD as an undergraduate at Caltech, although it was a particular incident involving an SVD question on a midterm exam in Charlie Van Loan's matrix computations class at Cornell University that helped inspire the original version of this article. As an exercise (which is related to the solution that he submitted for that exam), he encourages diligent readers to look up the many existing backronyms for SVD.

Mathematical Institute, University of Oxford, Oxford, OX1 3LB, United Kingdom
porterm@maths.ox.ac.uk