10. ———, Slinky models, *Phys. Teach.* **40** (2002) 324; available at http://dx.doi.org/10.1119/1.1511579
11. M. Sawicki, Static elongation of a suspended Slinky™, *Phys. Teach.* **40** (2002) 276–278; available at http://dx.doi.org/10.1119/1.1516379
12. T. P. Toepker, Center of mass of a suspended Slinky: an experiment, *Phys. Teach.* **42** (2004) 16–17; available at http://dx.doi.org/10.1119/1.1639962

An Intuitive Proof of the Singular Value Decomposition of a Matrix

Keith J. Coates (kcoates@drury.edu) Drury University, Springfield MO

The Singular Value Decomposition (SVD) is very well known. We provide an intuitive proof for real matrices. The following statement is adapted from [**1**]:

Singular Value Decomposition *Let the $m \times n$ real matrix A have rank k. Then there exist numbers $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_k > 0$, called the singular values of A, an $m \times m$ orthogonal matrix U and an $n \times n$ orthogonal matrix V such that $A = U\Sigma V^T$, where Σ is the $m \times n$ matrix*

$$\Sigma = \begin{pmatrix} D & 0 \\ 0 & 0 \end{pmatrix},$$

with D the $k \times k$ diagonal matrix with diagonal element $d_{ii} = \sigma_i$ for $1 \leq i \leq k$.

The matrices V and Σ are constructed inductively as follows. With v_i denoting the ith column of V, we let v_1 be a unit vector in $\mathcal{R}^n$ such that $\|Av_1\| = \|A\|$, and for $2 \leq i \leq k$, we let v_i be a unit vector in $\mathcal{R}^n$ such that

$$\|Av_i\| = \max\left\{\|Av\| : v \in \mathcal{B}^n \cap (\text{span}\{v_1, v_2, v_3, \ldots v_{i-1}\})^{\perp}\right\},$$

where $\mathcal{B}^n$ is the unit ball of $\mathcal{R}^n$ and $\|\cdot\|$ is the Euclidean norm or its induced matrix norm. The existence of each v_i follows from the continuity of the norm and the compactness of closed subsets of $\mathcal{B}^n$. The set $\{v_1, v_2, v_3 \ldots, v_k\}$ is obviously orthonormal. Setting $\sigma_i = \|Av_i\|$, we have $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_k \geq 0$. If $\sigma_i = 0$ for some i, then $\sigma_j = 0$ for all $j \geq i$ and the range of A is spanned by $\{Av_1, Av_2, Av_3, \ldots, Av_{i-1}\}$, contradicting the fact that A is of rank k. Thus $\sigma_i > 0$ for $1 \leq i \leq k$. This constructs D (and hence Σ), as well as the first k columns of V.

We define u_i, the ith column of U, to be $\frac{Av_i}{\sigma_i}$ for $1 \leq i \leq k$. Then $\|u_i\| = 1$. The key question is, why does this inductive construction lead to the pairwise orthogonality of $\{u_1, u_2, u_3, \ldots, u_k\}$? The remaining $n - k$ columns of V and $m - k$ columns of U are easily constructed once we establish orthogonality. We use an approach that is quite simple, based on the trigonometric limit

$$\lim_{\theta \to 0^+} \frac{1 - \cos\theta}{\sin\theta} = 0, \tag{1}$$

which in turn implies that for any $\epsilon > 0$, there exists $\delta > 0$ such that

$$\cos\theta + \epsilon \sin\theta > 1, \quad \text{for } 0 < \theta < \delta. \tag{2}$$

http://dx.doi.org/10.4169/college.math.j.42.5.394
MSC: 15A18, 15A23

Since each induction step is a simple modification of the first, we show only the first step. That is, we start with $v_1 \in \mathcal{B}^n$ such that $||Av_1|| = ||A||$ and show that for any w nonzero in $\mathcal{R}^n$ with $w \perp v_1$, we have $Aw \perp Av_1$. This implies that $u_i \perp u_1$ for $2 \leq i \leq k$.

Let $w \in \mathcal{R}^n$ satisfy $w \perp v_1$, $\|w\| = 1$, and $Aw \neq 0$. If such a w does not exist, then $k = 1$ and the induction is finished. It is a basic result from linear algebra that there is a real number ϵ and a vector $u \in \left(\text{span}\{Av_1\}\right)^{\perp} \subseteq \mathcal{R}^m$ such that $Aw = \epsilon Av_1 + u$. By replacing w with $-w$, if necessary, we may assume $\epsilon \geq 0$. To show that $Aw \perp Av_1$, we will show that $\epsilon = 0$.

With θ acute and positive, let y_θ be the unit vector $(\cos\theta)v_1 + (\sin\theta)w \in \mathcal{R}^n$. Then

$$\begin{aligned} Ay_\theta &= (\cos\theta)Av_1 + (\sin\theta)Aw \\ &= (\cos\theta)Av_1 + (\sin\theta)(\epsilon Av_1 + u) \\ &= (\cos\theta + \epsilon\sin\theta)Av_1 + (\sin\theta)u. \end{aligned}$$

The norm of Ay_θ is at least as large as its component in the direction of Av_1: $\|Ay_\theta\| \geq (\cos\theta + \epsilon\sin\theta)\|Av_1\|$. But if ϵ is nonzero, by (2) we can choose $\theta > 0$ small enough so that $(\cos\theta + \epsilon\sin\theta)\|Av_1\| > \|Av_1\|$. Then y_θ would be a unit vector and $\|Ay_\theta\| > \|Av_1\| = \|A\|$, a contradiction. We conclude that $\epsilon = 0$ and $Aw \perp Av_1$.

To complete the proof of the SVD, we point out that if $v \in \left(\text{span}\{v_1, v_2, v_3, \ldots, v_k\}\right)^{\perp}$, then the above argument implies that $Av \in \left(\text{span}\{u_1, u_2, u_3, \ldots, u_k\}\right)^{\perp}$. Since A has rank k, this implies that $Av = 0$. Otherwise, $\{u_1, u_2, u_3, \ldots, u_k, Av\}$ would be a linearly independent set in the range of A with cardinality $k + 1$. Thus we may choose the remaining $n - k$ columns of V in any fashion that extends $\{v_1, v_2, v_3, \ldots, v_k\}$ to an orthonormal basis of $\mathcal{R}^n$, and likewise choose the remaining $m - k$ columns of U to extend $\{u_1, u_2, u_3, \ldots, u_k\}$ to an orthonormal basis of $\mathcal{R}^m$. This completes the construction of the SVD of the matrix A.

Figure 1 is a visual depiction of the proof. Starting with the orthogonal unit vectors v_1 and w and applying A to the unit vector $y_\theta = (\cos\theta)v_1 + (\sin\theta)w$ in $\mathcal{R}^n$, we have $Ay_\theta = (\cos\theta)Av_1 + (\sin\theta)Aw$. If Aw were orthogonal to Av_1, then we could regard Ay_θ as having "horizontal" vector component $(\cos\theta)Av_1$ and "vertical" component $(\sin\theta)Aw$. But, under our assumption that Aw is *not* orthogonal to Av_1, we have $Aw = \epsilon Av_1 + u$, so that $(\sin\theta)Aw = (\epsilon\sin\theta)Av_1 + (\sin\theta)u$, with $\epsilon > 0$ and $u \in \left(\text{span}\{Av_1\}\right)^{\perp}$. The vector $(\sin\theta)u$ is now the vertical vector component of Ay_θ and $(\epsilon\sin\theta)Av_1$ gets added to $(\cos\theta)Av_1$ to give us the horizontal component of Ay_θ. In the figure we see that as long as θ is small enough and nonzero, this extra piece added in the horizontal direction is enough to make the horizontal component of Ay_θ longer than the length of Av_1, giving us our contradiction.

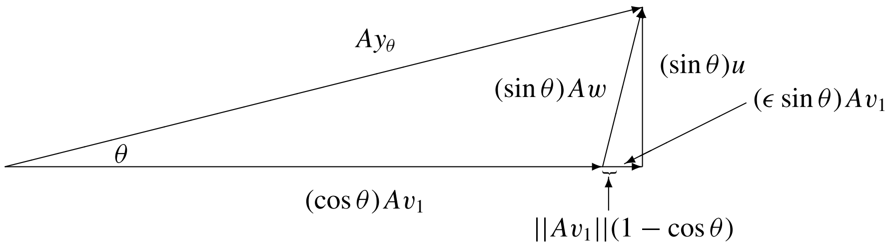

Figure 1.

Summary. Using a simple trigonometric limit, we provide an intuitive geometric proof of the Singular Value Decomposition of an arbitrary matrix.

Acknowledgment. The author thanks Scott J. Simmons, a CMJ Classroom Capsule editor, the CMJ senior editor, and an anonymous referee for many helpful suggestions.

References

1. B. Noble and J. W. Daniel, *Applied Linear Algebra*, Prentice-Hall, Englewood Cliffs NJ, 1977.

Discretization vs. Rounding Error in Euler's Method

Carlos F. Borges (borges@nps.edu), Department of Applied Mathematics, Naval Postgraduate School, Monterey CA

Euler's method for solving initial value problems is a good vehicle for observing the relationship between discretization error and rounding error in numerical computation. As we reduce stepsize, in order to decrease *discretization error*, we necessarily increase the number of steps and introduce additional *rounding error*. The problem is common and can be quite troublesome. We will examine a simple device that can help delay the onset of this problem.

Meet the Eulers Consider the problem of solving the ordinary differential equation $\dot{x} = f(x, t)$ on the interval $[0, T]$, subject to the initial condition $x(0) = x_0$. A common technique for solving such a problem numerically is Euler's method. The method starts by selecting a positive integer n and discretizing the time axis into a set of lattice points $t_k = kh$ for $k = 0, 1, 2, \ldots, n$, where $h = T/n$ is called the *step size*. Provided that that $x(t)$ is sufficiently differentiable, we have, by Taylor's formula,

$$x(t_k + h) = x(t_k) + h\dot{x}(t_k) + O(h^2).$$

Replacing $t_k + h$ by t_{k+1} and $\dot{x}(t_k)$ by $f(x(t_k), t_k)$ yields

$$x(t_{k+1}) = x(t_k) + hf(x(t_k), t_k) + O(h^2).$$

Ignoring the *single-step discretization error* $O(h^2)$ leads to the Euler approximation. Euler's method is the repeated use of this approximation across the entire lattice of discretized time points. In particular, we set $\psi_0 = x_0$ and use the recurrence

$$\psi_{k+1} = \psi_k + hf(\psi_k, t_k) \tag{1}$$

for $k = 0, 1, \ldots, n-1$. If the above scheme is carried out in exact arithmetic, then we call ψ_n the *Euler solution* and the difference $x(T) - \psi_n$ the *discretization error* for step-size h.

A less common derivation is to transform the initial value problem to the definite integral:

$$x(T) = x_0 + \int_0^T f(x, t)dt,$$

http://dx.doi.org/10.4169/college.math.j.42.5.396
MSC: 65-01, 65-04, 65L05