


This article was downloaded by: [Colorado School of Mines], [Stephen Pankavich]

On: 19 June 2015, At: 09:24

Publisher: Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954

Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



PRIMUS: Problems, Resources, and Issues in Mathematics Undergraduate Studies

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/upri20>

Principal Component Analysis: Resources for an Essential Application of Linear Algebra

Stephen Pankavich & Rebecca Swanson

Accepted author version posted online: 20 Dec 2014. Published online: 20 Dec 2015.



CrossMark

[Click for updates](#)

To cite this article: Stephen Pankavich & Rebecca Swanson (2015) Principal Component Analysis: Resources for an Essential Application of Linear Algebra, PRIMUS: Problems, Resources, and Issues in Mathematics Undergraduate Studies, 25:5, 400-420, DOI: [10.1080/10511970.2014.993446](https://doi.org/10.1080/10511970.2014.993446)

To link to this article: <http://dx.doi.org/10.1080/10511970.2014.993446>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with

primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

Principal Component Analysis: Resources for an Essential Application of Linear Algebra

Stephen Pankavich and Rebecca Swanson

Abstract: Principal Component Analysis (PCA) is a highly useful topic within an introductory Linear Algebra course, especially since it can be used to incorporate a number of applied projects. This method represents an essential application and extension of the Spectral Theorem and is commonly used within a variety of fields, including statistics, neuroscience, and image compression. We present a synopsis of PCA and include a number of examples that can be used within upper-level mathematics courses to engage undergraduate students while introducing them to one of the most widely used applications of linear algebra.

Keywords: Linear algebra, principal component analysis, data analysis, neuroscience, image compression, *MATLAB*.

1. INTRODUCTION

Although it is often the first proof-based course most students experience, linear algebra is also an important topic that is used to solve problems in a wide variety of fields. This is why, at many universities, the course is a requirement for a large number of degree programs in engineering and the natural sciences. However, since many colleges, especially those with a technical focus, possess significantly more of these majors than those within mathematics, the majority of enrolled students often find difficulty with the abstract nature of the subject, especially when the questions “How does this apply to my major?” and “What are the real-world applications?” can largely go unanswered.

Address correspondence to Stephen Pankavich, Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, CO 80403, USA. E-mail: pankavic@mines.edu

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/upri

One valuable application that can be included in a first Linear Algebra course is the method of Principal Component Analysis (PCA). PCA is a sensible choice for an applied example to which linear algebra is crucial because it arises in so many different contexts, as we will demonstrate within subsequent sections. The method arises in countless disciplines, including but not limited to statistics [4], electrical engineering [10], genetics [9], neuroscience [8], facial recognition [15], control theory [10], and mechanical and systems engineering [1]. Perhaps the largest obstacle to including PCA within a first semester Linear Algebra course is that it is often performed using the Singular Value Decomposition (SVD) approach, which is a topic usually reserved for either a second course in the subject, a course in computational linear algebra, or a graduate linear algebra course (see [3, 5, 11] for more information on SVD). Of course, many universities do not offer these follow-on courses, and hence students are not exposed to PCA, even though it is such a widely utilized tool. However, we note that PCA can be both understood and performed without knowledge of the SVD approach, using a topic that is usually included within a first-semester course, namely the Spectral Theorem. We will elaborate further on how this is accomplished later on. In addition, although PCA is sometimes included as an applied topic within linear algebra texts, most notably [6] and [13], it is often overlooked as an important application within other sources. Even when discussed in these texts, it is often difficult to describe the utility or impact of the method because examples within textbooks are typically computed by hand, whereas real-world implementation is always performed using a computer.

To address these difficulties, we have created resources, both theoretical and computational, for instructional use within the classroom of a first course in linear algebra. These resources require a computer in the classroom, but will not require students to understand numerical aspects of linear algebra. We also provide the code that generates the data sets, numerics, and images used within our examples. Instructors who are comfortable with *MATLAB* can alter our code to make the presentation more interactive; however, those with less computational interest or experience need not do so in order to generate all of the material we present. Although we use *MATLAB* to perform computations, the principles and algorithms can be applied using other (perhaps open-source) software, such as *Octave*. For instructors without access to *MATLAB*, we also provide a number of programs on the first author's website: <http://inside.mines.edu/~pankavic/PCACode/Octave> that are similar to those in the Appendix but written using *Octave*.

These computational tools arise from a first-semester Linear Algebra course conducted at the Colorado School of Mines, and have had a strong impact on the interest of non-majors and, subsequently, the increase in mathematics minors at our institution. Both authors have implemented portions of this material in their classrooms and have received positive feedback from students regarding their interest in PCA applications, especially concerning its

use in image compression and statistics. In addition, the inclusion of PCA has benefitted the mathematical breadth and depth of our applied majors and increased their exposure to new and interesting applications of linear algebra. Hence, we believe the resources included herein can be used to introduce students to a variety of real-world applications of the subject without the need for a course in scientific computing.

We also mention that the implementation of these materials is fairly robust. Instead of a Linear Algebra course, they could be included within an introductory or second course in statistics for students who have already taken linear algebra, or within a scientific computing course, assuming students possess the necessary background. In general, this article provides a number of teaching resources and examples to engage students in learning about this essential application of linear algebra.

2. BACKGROUND

We begin by recalling a crucial theorem - one of the most important results in all of linear algebra - on which we base our discussion of PCA. Before stating the result, we note that a stronger version of the theorem extends to a more general class of matrices over the complex numbers.

Theorem. (*Spectral Theorem [6]*) *Let $n \in \mathbb{N}$ be given and A be an $n \times n$ matrix of real numbers with A^T denoting its transpose. Then, A is symmetric (i.e., $A^T = A$) if and only if A is orthogonally diagonalizable. In this case, A possesses n real eigenvalues, counting multiplicities.*

This is an extremely powerful result and precisely guarantees for such A the existence of $\lambda_k \in \mathbb{R}$ and orthonormal column vectors $\mathbf{v}_k \in \mathbb{R}^n$ for every $k = 1, \dots, n$ such that

$$A = \lambda_1 \mathbf{v}_1 \mathbf{v}_1^T + \cdots + \lambda_n \mathbf{v}_n \mathbf{v}_n^T \quad (1)$$

arising from the orthogonal diagonalization. Because the \mathbf{v}_k vectors are orthonormal, each associated matrix $\mathbf{v}_k \mathbf{v}_k^T$ is orthogonal, and thus A can be decomposed into a sum wherein each term is an eigenvalue multiplied by a rank-one matrix generated by a unit vector. Hence, the eigenvalues alone determine the magnitude of each term in the sum, whereas the eigenvectors determine the directions. These eigenvectors are called *principal components* or *principal directions* and we will expand upon this further in the next section. We begin the discussion of the specifics of PCA with an introductory example.

2.1. Introductory Height and Weight Problem

Consider a study in which we want to determine whether or not the heights and weights of a group of individuals are correlated. That is, we want to know whether the known value of a person’s height seems to dictate whether they tend to be heavier or lighter, and thus influences their weight. Assume we are given data for 30 specific people, displayed within [Table 1](#) and [Figure 1](#). For this example, our data set originates from a commonly available study [12]. In practice and for large enough class sizes, the data can be obtained in the classroom, directly from students. For instance, during the class period before implementation of this example, the instructor can collect anonymous information regarding student height and weight, compile and store this data, and integrate it into the project for the following class period. As an alternative, this problem could serve as a project for a group of students within the course, wherein the group collects the data, performs PCA using the code provided in the Appendix, interprets the results, and presents their findings in a brief class seminar. For those with smaller class sizes or concerns regarding anonymity, data can be taken from our source [12] that is readily available on the Internet.

Table 1. Heights (in.) and weights (lbs) for 30 young adults [12]

Person	1	2	3	4	5	6
Height	67.78	73.52	71.40	70.22	69.79	70.70
Weight	132.99	176.49	173.03	162.34	164.30	143.30
Person	7	8	9	10	11	12
Height	71.80	72.01	69.90	68.78	68.49	69.62
Weight	161.49	166.46	142.37	150.67	147.45	144.14
Person	13	14	15	16	17	18
Height	70.30	69.12	70.28	73.09	68.46	70.65
Weight	155.61	142.46	146.09	175.00	149.50	162.97
Person	19	20	21	22	23	24
Height	73.23	69.13	69.83	70.88	65.48	70.42
Weight	177.90	144.04	161.28	163.54	126.90	149.50
Person	25	26	27	28	29	30
Height	69.63	69.21	72.84	69.49	68.53	67.44
Weight	161.85	149.72	172.42	151.55	138.33	133.89




Figure 1. Plot of height/weight datapoints.

Since the question of interest is whether the two measured variables, height and weight, seem to change together, the relevant quantity to consider is the covariance of the two characteristics within the data set. This can be formed in the following way. First, the data is stored in X , a 2×30 matrix. Then, the entries are used to compute the mean in each row, which will be used to center or “mean-subtract” the data. This latter step is essential, as many of the results concerning PCA are only valid upon centering the data at the origin. Computing the means of our measurements (Table 1), we find

$$\boldsymbol{\mu} = \begin{bmatrix} 70.06 \\ 154.25 \end{bmatrix}.$$

Using x_{ij} , the entries of the data matrix X , the associated 2×2 covariance matrix S is constructed with entries

$$s_{ik} = \frac{1}{30-1} \sum_{j=1}^{30} (x_{ij} - \mu_i)(x_{kj} - \mu_k),$$

so that

$$S = \begin{bmatrix} 3.26 & 21.72 \\ 21.72 & 188.96 \end{bmatrix}.$$

Notice that this matrix is necessarily symmetric, so using the Spectral Theorem it can be orthogonally diagonalized. Upon computing the eigenvalues and eigenvectors of S , we find $\lambda_1 = 191.46$, $\lambda_2 = 0.76$, and

$$\mathbf{v}_1 = \begin{bmatrix} 0.11 \\ 0.99 \end{bmatrix} \quad \mathbf{v}_2 = \begin{bmatrix} -0.99 \\ 0.11 \end{bmatrix}.$$

Here, \mathbf{v}_1 and \mathbf{v}_2 are the *principal components* of the covariance matrix S generated by the data matrix X , as previously described. Thus, we see from (1) that

$$S = \lambda_1 \mathbf{v}_1 \mathbf{v}_1^T + \lambda_2 \mathbf{v}_2 \mathbf{v}_2^T,$$

and because the difference in eigenvalues is so large, it appears that the first term is responsible for most of the information encapsulated within S . Regardless, we can re-express the given data in the new orthonormal basis generated by \mathbf{v}_1 and \mathbf{v}_2 by computing the coordinates $P^T X$ where

$$P = \begin{bmatrix} 0.11 & -0.99 \\ 0.99 & 0.11 \end{bmatrix}$$

is the orthogonal matrix whose columns are \mathbf{v}_1 and \mathbf{v}_2 . In fact, we could left multiply the data matrix by each component separately, namely $\mathbf{v}_1^T X$ and $\mathbf{v}_2^T X$, to project the data onto each principal direction (Figure 2). Hence, the data can be separated into projections along \mathbf{v}_1 and \mathbf{v}_2 , respectively. We see from looking at the scales in Figure 2 that the heights and weights along \mathbf{v}_2 are significantly




Figure 2. Height/weight data projected onto the principal components (left - \mathbf{v}_1 ; right - \mathbf{v}_2). By the Spectral Theorem, the data represented in Figure 1 is exactly the sum of the projections onto these two components.

less than those along \mathbf{v}_1 , which tells us that the majority of the information contained within X lies along \mathbf{v}_1 . Computing the slope of the line in the direction of \mathbf{v}_1 and choosing a point thru which it passes, we can represent it by

$$y - 154.25 = 9(x - 70.06),$$

where x represents the height of a given individual and y is their corresponding weight. Hence, we see that height and weight appear to be strongly correlated, and PCA has determined the direction with optimal correlation between the variables.

The PCA for this example took a small set of data and identified a new orthonormal basis in which to re-express it. In two dimensions the data are effectively rotated to lie along the line of best fit (Figure 3), with the second principal direction merely representing the associated unit orthogonal complement of the first. This mirrors the general aim of PCA; to obtain a new orthonormal basis that organizes the data optimally, in the sense that the variance contained within the vectors is maximized along successive principal component(s).




Figure 3. Biplot of height/weight data with two principal components. The blue height and weight vectors are displayed as linear combinations of the principal components. Note that these principal components effectively rotate the height and weight data in the plane.

2.2. Summary of PCA

In short, PCA can be performed to compute an optimal, ordered orthonormal basis of a given set of vectors, or data set, in the following steps.

1. Gather n samples of m -dimensional data, i.e. vectors $\mathbf{d}_1, \dots, \mathbf{d}_n \in \mathbb{R}^m$ stored in the $m \times n$ matrix X with columns $\mathbf{d}_1, \dots, \mathbf{d}_n$, so that x_{ij} represents the i th entry of the j th sample vector, and compute the mean vector (in \mathbb{R}^m)

$$\boldsymbol{\mu} = \frac{1}{n} \sum_{k=1}^n \mathbf{d}_k.$$

2. Build the corresponding mean-centered data matrix B with columns given by $\mathbf{d}_j - \boldsymbol{\mu}$ so that the entries are

$$b_{ij} = x_{ij} - \mu_i$$

for every $i = 1, \dots, m$ and $j = 1, \dots, n$.

3. Use B to compute the symmetric, $m \times m$ covariance matrix

$$S = \frac{1}{n-1} BB^T.$$

4. Find the eigenvalues $\lambda_1, \dots, \lambda_m$ of S (arranged in decreasing order including multiplicity) and an orthonormal set of corresponding eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_m$. These create a new basis for \mathbb{R}^m in which the data can be expressed.
5. Finally, the data is represented in terms of the new basis vectors $\mathbf{v}_1, \dots, \mathbf{v}_m$ using the coordinates $\mathbf{y}_1 = \mathbf{v}_1^T X, \dots, \mathbf{y}_m = \mathbf{v}_m^T X$. This can also be represented as the matrix $Y = P^T X$ where P is the matrix with columns $\mathbf{v}_1, \dots, \mathbf{v}_m$. Should we wish to convert the data back to the original basis, we merely utilize the orthogonality of P and compute PY to find X .

In the final sections, we will extend our introductory example while presenting additional applications in which PCA appears prominently and can be implemented within a classroom environment. For each of the subsequent examples, an instructor could integrate the content into handouts or group worksheets, assign group projects with brief presentations during class, or merely use the given materials to present the information in an interactive lecture format.

3. APPLICATION: DATA ANALYSIS

In the previous section, we developed a method for PCA that determined a basis with maximal variance. Notice that the first component really encapsulated the majority of the information embedded within the data. Since the eigenvalues can be ordered, we might also be able to truncate the sum in (1) to reduce the amount of stored data. For instance, in our previous example, we might only keep the first principal component since the data can be mostly explained just by knowing this characteristic, rather than every height and weight. In this case, each data point would then be represented by its projection onto the first principal component. Upon performing the final step, we might also interpret the results: Are a small number of the eigenvalues λ_k much less (perhaps by an order of magnitude) than the others? If so, this indicates a reduction in the dimension of the data is possible without losing much information; however, if this does not occur then the dimension of the data may not be easily reduced in this way.

Suppose that in addition to computing the components, we were to truncate the new basis matrix, P , so that we keep only the first r columns with $r < m$. We would thus have a matrix $\tilde{P} \in \mathbb{R}^{m \times r}$, and this would give rise to the $r \times n$ matrix $\tilde{Y} = \tilde{P}^T X$ containing a truncation of the data represented by the first r principal components. From this we could also create the $m \times n$ matrix $\tilde{X} = \tilde{P}\tilde{P}^T X$, which represents the projection of X onto the first r principal directions. This reduced representation of the data would then possess less information than X , but retain the most information when compared to any other matrix of the same rank. In fact, an error estimate is obtained from the Spectral Theorem - namely, the amount of information retained is given by the *spectral ratio* of the associated covariance matrix

$$\sigma^2 := \frac{\sum_{k=1}^r \lambda_k}{\sum_{k=1}^n \lambda_k}. \quad (2)$$

Thus, in our first example, we can keep only the first component of each data point (a 1×30 matrix) rather than the full data set (2×30 matrix) and still retain 99% of the information contained within because

$$\sigma^2 = \frac{191.46}{191.46 + 0.76} > 0.99.$$

In situations where the dimension of the input vector is large, but the components of the vectors are highly correlated, it is beneficial to reduce the dimension of the data matrix using PCA. This has three effects: (i) it orthogonalizes the basis vectors (so that they are uncorrelated); (ii) it orders the resulting orthogonal components so that those with the largest variance appear

first; and (iii) it eliminates dimensions that contribute the least to the variation in the data set.

We now extend our introductory example to larger data sets with many other characteristics. For our first example, we will utilize the built-in *MATLAB* data set `fisheriris.mat`. This famous collection of data arises from Fisher’s 1936 paper [2] that describes 50 different samples of four characteristics from each of three species of iris. Hence, the data set contains 150 points and four variables: sepal length, sepal width, petal length, and petal width. With the 4×150 data matrix loaded, we perform the steps outlined within Section 2.2 and compute the first two principal components, i.e., those corresponding to the largest eigenvalues $\lambda_1 = 4.23$ and $\lambda_2 = 0.24$. The others $\lambda_3 = 0.08$ and $\lambda_4 = 0.02$ are omitted.

It is difficult to visualize even the first two principal components in this example because they are vectors in \mathbb{R}^4 , but we can list them as in Table 2.

Alternatively, we can visualize each data point projected onto the first two principal components as in Figure 4(a). This figure is actually a biplot, containing both the projected data points and the proportion of each characteristic that accounts for the respective principal component. For instance, because the petal length accounts for a large proportion of the first principal component, it points in nearly the same direction as the x -axis within the biplot. Similarly, sepal width and sepal length seem to account for a large amount of the second principal component.

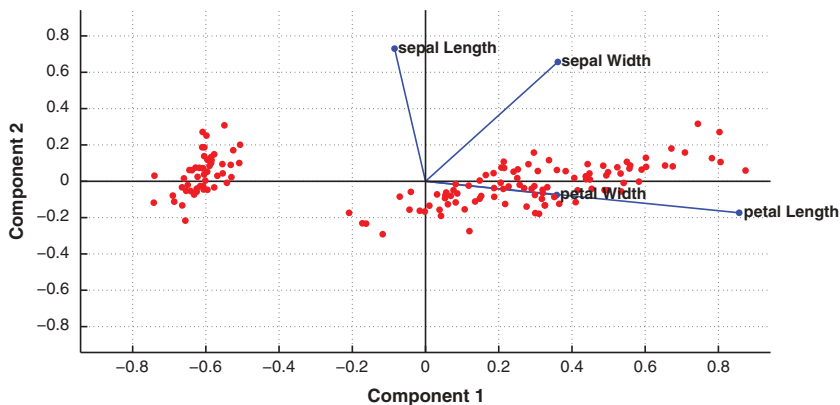
Notice in Figure 4(a) that most of the data lie along the first component, which is mostly determined by the length of the iris petals. Additionally, petal length and petal width appear to be strongly correlated because they point in nearly identical directions within the new basis. This can be seen in the enlarged portion of the biplot shown in Figure 4(b). Contrastingly, sepal width and sepal length seem only mildly correlated and neither appears to be correlated to the pedal characteristics since these vectors point in somewhat orthogonal directions.

Of course, the dimension and complexity of this example can be increased by using a larger set of data, such as another built-in *MATLAB* sample called

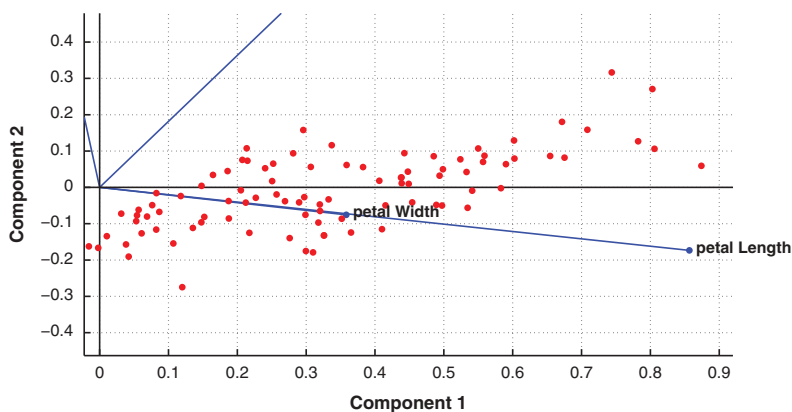
Table 2. The first two principal components

Characteristic	PC1	PC2
sepal length	0.3614	0.6566
sepal width	-0.0845	0.7302
petal length	0.8567	-0.1734
petal width	0.3583	-0.0755

Downloaded by [Colorado School of Mines], [Stephen Pankavich] at 09:24 19 June 2015



(a)



(b)

Figure 4. Biplots of Fisher's iris data projected onto $r = 2$ principal components with $\sigma^2 = 97.77\%$. Note that (b) is an enlarged portion of (a). Also note that petal Width and petal Length point in nearly identical directions.

`cities.mat`. For completeness, we have included biplots (Figure 5) of the first few principal components of the `cities` set, which contains $m = 9$ different attributes (i.e., climate, crime, education, etc.) for $n = 329$ cities, again using the code contained in the Appendix. As a final observation, we note that due to the separation between eigenvalues of the covariance matrix (sometimes called the “spectral gap”), the majority of data points lie along the first principal component in Figure 5(a) and within the plane generated by the first two principal components in Figure 5(b). This indicates that the majority of the variance




Figure 5. Biplots of cities data, both with $\sigma^2 \approx 94.43\%$: (a) $r = 2$ components and (b) $r = 3$ components.

within the data exists in these two directions, and hence one might safely eliminate the other dimensions, which contribute the least to the variation in the data set.

4. APPLICATION: NEUROSCIENCE

For another field in which PCA is quite useful, we turn to neuroscience. In electrophysiological recordings of neural activity, one electrode typically records the action potentials (or spikes) of several neurons. Before one can use the recorded spikes to study the coding of information in the brain, they must first be associated with the neuron(s) from which the signal arose. This is often accomplished by a procedure called *spike sorting*, which can be accomplished because the recorded spikes of each neuron often have characteristic shapes [7]. For example, Figure 6 shows three different shapes of action potentials recorded with an electrode. These potentials are plotted by connecting measurements at 64 different points in time using linear interpolation. Due to distinctions in their peaks and oscillations, the spikes in Figure 6 are due to three different neurons, and PCA can be used to identify the principal variations of these spike shapes.

Let us consider a data matrix X of 9195 recorded spikes, sampled at 64 different points in time as in the previous figure so that $X \in \mathbb{R}^{64 \times 9195}$. For this example, we use real human data that is freely available online [7]; however, note that this data is not directly connected to the examples of action potentials in Figure 6. Using PCA we compute the eigenvalues of the covariance matrix to find that the first two ($\lambda_1 = 2.9 \times 10^4$ and $\lambda_2 = 4.7 \times 10^3$) account for 83% of the total information. Of course, additional components can be included to increase this percentage. A plot of the data projected along the




Figure 6. Examples of action potentials.

first two principal components is given in [Figure 7](#). We notice that two distinct clusters form within the data, and thus it appears that the spikes are formed from two different neurons. If we denote the first two principal vectors by \mathbf{v}_1 and \mathbf{v}_2 , then Cluster #1 (right) and Cluster #2 (left) appear to center around $100\mathbf{v}_1 + 25\mathbf{v}_2$ and $-275\mathbf{v}_1$, respectively. The spike shapes corresponding to these vectors are displayed in [Figure 8](#) and represent averaged activity from the two neurons. Hence, we see that PCA can be used both to determine the degree




Figure 7. Spike data projected onto principal components.




Figure 8. Representative action potentials for clusters in projected data.

of correlation among certain characteristics and to identify clustering patterns within data. Additionally, it provides us with a lower-dimensional picture of high-dimensional data sets, and this is often very useful when attempting to visualize high-dimensional data. With the action potentials determined, they can be associated to specific neurons and analyzed within neuroscience studies.

5. APPLICATION: IMAGE COMPRESSION

Another important application of PCA is image compression. Because images are stored as large matrices with real entries, one can reduce their storage requirements by keeping only the essential portions of the image [6]. Of course, information (in this case, fine-grained detail of the image) is naturally lost in this process, but it is done so in an optimal manner, so as to maintain the most essential characteristics.

In this section we detail a specific example for the use of PCA to compress an image. Since the effects of keeping a lower-dimensional projection of the image will be visually clear, this particular example is a great candidate for an interactive, in-class activity. More specifically, one can provide students with the code given in the Appendix and ask them to determine the number of principal components r that they must preserve in order to visually identify the




Figure 10. The first 35 eigenvalues of the covariance matrix S generated from X in the image compression example and arranged in decreasing order.

the truncation of the sum of principal components in \tilde{X} should still contain a large amount of the total information of the original image X . In Figure 11, we represent \tilde{X} for four choices of r (i.e., the number of principal components used), and the associated spectral ratio, σ^2 , retained by those reduced descriptions is also listed. Notice that the detail of the image improves as r is increased, and that a fairly suitable representation can be obtained with around 90 components rather than the full 648 vector description. Hence, PCA has again served the useful purpose of reducing the dimension of the original data set while preserving its most essential features.

6. CONCLUSIONS

The aim in writing this article is to generate and present clear, concise resources for the instruction of PCA within an introductory Linear Algebra course. Students should be led to understand that PCA is a powerful, useful tool that is utilized in data analysis and throughout the sciences. The visual components of the image compression examples in particular are designed to display the utility of PCA and linear algebra, in general. Within a more advanced setting (perhaps a second-semester course in linear algebra), new material regarding the SVD approach can also be included to introduce additional details concerning the implementation of PCA and provide an elegant structure for performing the method [5, 14]. That being said, the SVD can often be a technical and time-consuming topic that one first encounters

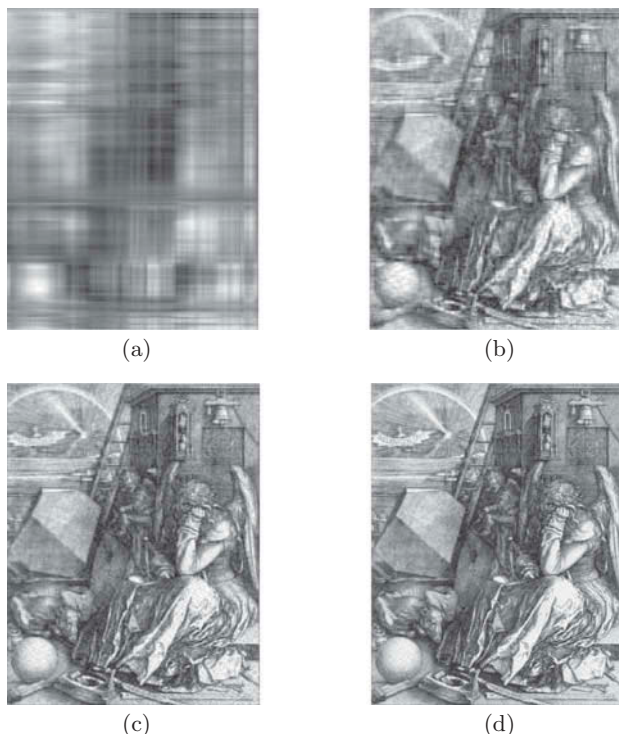


Figure 11. The Dürer image with varying numbers of principal components: (a) $r = 3$, $\sigma^2 = 88.89\%$; (b) $r = 30$, $\sigma^2 = 99.61\%$; (c) $r = 60$, $\sigma^2 = 99.78\%$; and (d) $r = 90$, $\sigma^2 = 99.92\%$.

in a computational, honors, or graduate linear algebra course rather than within an introductory setting. In general, we believe that these instructional resources should be helpful in seamlessly integrating one of the most essential current applications of linear algebra, PCA, into a standard undergraduate course.

FUNDING

The authors would like to thank the College of Engineering and Computational Sciences at the Colorado School of Mines for partial support to develop and implement the educational tools included within this article. Additionally, the first author thanks the National Science Foundation for support under award DMS-1211667.

APPENDIX

MATLAB Code for Examples

Load height/weight, iris, cities, neuroscience, or Durer image data and call functions. Recall that the “HtWtdata” and “times_CSC4” data sets cannot be loaded directly from *MATLAB*, but can be obtained online using the aforementioned references.

```
1 %Initialize data for one of the five examples below
2
3 %%% For Height/Weight example, use
4 % load HtWtdata
5 % vbls = {'Height', 'Weight'};
6 % vbls=[];
7 % plottype = 1;
8
9 %%% For Iris data example, use
10 % load fisheriris
11 % data = meas';
12 % vbls = {'Sepal Width', 'Sepal Height', 'Petal ...
    Length', 'Petal Width'};
13 % plottype = 2;
14
15 %%% For Cities example, use
16 % load cities; data = ratings'; vbls = categories;
17 % plottype = 2;
18
19 %%% For Durer example, use
20 % load durer; data=double (X)
21 % fig1=figure ('Color',[1 1 1]);
22 % image (X), colormap (map), axis off, axis equal; % plot
23 % vbls=[] ; plottype = 3;
24
25 %%% For Neuroscience example, use
26 % load times_CSC4
27 % data = spikes';
28 % vbls=[];
29 % plottype = 4; % plottype = 5;
30
31 PC = 2; %Specify number of components to use
32 %Function calls for PCA and plotting
33 [X, V, Y] = PCAfun (data, PC);
34 PlotData (plottype, X, V, Y, PC, vbls, data);
```

Function to plot components, biplots, and images

```

1 function PlotData(num, X, V, Y, PC, vbls, data)
2
3 fig3 = figure('Color',[1 1 1]);
4 if num == 1           % Plot for Statistics example
5   plot(X(1,:), X(2,:), '.', 'MarkerSize', 30)
6   xlabel('Height (in)')
7   ylabel('Weight (lbs)')
8
9 elseif num == 2      % Biplot for Statistics example
10  Y = Y';
11  biplot(V, 'scores',Y(:,1:PC), 'MarkerSize', ...
12         20, 'varlabels',vbls)
13
14 elseif num == 3      % Plot for Image Compression example
15  image(X, colormap(map), axis off, axis equal; % ...
16         display results
17
18 elseif num == 4      % Plot for Neuroscience Cluster
19  plot(Y(1,:), Y(2,:), '.', 'MarkerSize', 20), axis([-400, ...
20         200, -200, 200])
21  xlabel('Component 1 ', 'FontSize', 10)
22  ylabel('Component 2 ', 'FontSize', 10)
23
24 elseif num == 5      % Plot for Action Potentials
25  time = linspace(-0.5, 0.5, 64);
26  mu = mean(data,2);
27  pot1 = mu + 110*V(:,1) - 25*V(:,2);
28  pot2 = mu - 260*V(:,1) - 10*V(:,2);
29  plot(time, pot1, time, pot2, '--r')
30  xlabel('Times (ms)', 'Interpreter', 'latex')
31  ylabel('Potential ($\mu$V)', 'Interpreter', 'latex')
32  legend('Cluster #1', 'Cluster #2')
33
34 end
35
36 figureHandle =(gcf);
37 %% make all text in the figure to size 16 and bold
38 set(findall (figureHandle, 'type', 'text'), 'fontSize',
39 16, 'fontWeight', 'bold')

```

Function to compute principal components from given data

```

1 function [Xtrunc, Vtrunc, Y] = PCAfun (data, PC)
2 % inputs load data and establish number of components to use
3
4 [m n] = size (data); % obtain dimensions of data matrix
5 mu = mean (data,2); % compute row mean
6 X = data-repmat (mu,1,n); % subtract row mean to obtain X
7 Z = 1/sqrt (n-1)*X'; % create matrix, Z
8 covar = Z'*Z; % covariance matrix of Z
9
10 % Extract eigenvalues and eigenvectors
11 [V, Lambda] = eigs (covar); % compute eigenvectors and ...
    eigenvalues
12 evals = diag (Lambda); % store eigenvalues in vector form
13 fig2 = figure ('Color',[1 1 1]);
14 bar (evals); % bar plot of eigenvalues (variances)
15
16 % compute percentage of variance included
17 totvals = sum(evals); %total sum of eigenvalues
18 runtot = cumsum(evals); %cumulative sum of ...
    eigenvalues - largest to smallest
19 percentinc = runtot(PC)/totvals % percentage included
20
21 % Extract principal components
22 Vtrunc = V(:,1:PC); % extract first PC vectors in new basis
23 Y = Vtrunc' * X; % project data onto new basis
24 Xtrunc = Vtrunc*Y; % convert back to original basis
25 Xtrunc = Xtrunc+repmat(mu,1,n); % add the row means back

```

REFERENCES

1. Feeny, B. F. and R. Kappagantu. 1998. On the physical interpretation of proper orthogonal modes in vibrations. *Journal of Sound and Vibration*. 211(4): 607–616.
2. Fisher, R. A. 1936. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*. 7(2): 179–188.
3. James, D., M. Lachance, and J. Remski. 2011. Singular vectors' subtle secrets. *The College Mathematics Journal*. 42(2): 86–95.
4. Jolliffe, I. T. 2002. *Principal Component Analysis, Second Edition*. New York: Springer.
5. Kalman, D. 1996. A singularly valuable decomposition: The SVD of a matrix. *The College Mathematics Journal*. 27: 2–23.
6. Lay, D. 2012. *Linear Algebra and its Applications, Fourth Edition*. Boston, MA: Pearson.
7. NeuroEngineering Laboratory, University of Leicester. <http://www2.le.ac.uk/departments/engineering/research/bioengineering/neuroengineering-lab/spike-sorting>. Last accessed September 2014.

8. Peyrache, A., K. Benchenane, M. Khamassi, S. Wiener, and F. Battaglia. 2010. Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution. *Journal of Computational Neuroscience*. 29(1–2): 309–325.
9. Raychaudhuri, S., J. Stuart, and R. Altman. 2000. Principal components analysis to summarize microarray experiments: Application to sporulation times series. In *Proceedings of the Pacific Symposium on Biocomputing* pp. 455–466.
10. Sanguansat, P. (Ed.). 2012. *Principal Component Analysis - Engineering Applications*. Rijeka, Croatia: InTech.
11. Shlens, J. A tutorial on principal component analysis. <http://arxiv.org/abs/1404.1100>. Last accessed September 2014.
12. SOCR. Data on human heights and weights. http://wiki.stat.ucla.edu/socr/index.php/SOCR_Data_Dinov_020108_HeightsWeights. Last accessed September 2014.
13. Strang, G. 2005. *Linear Algebra and its applications, Fourth Edition*. Boston, MA: Cengage Learning.
14. Trefethen, L. and D. Bau. 1997. *Numerical Linear Algebra*. Philadelphia, PA: SIAM.
15. Turk, M. and A. Pentland. 1991. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*. 3(1): 71–86.

BIOGRAPHICAL SKETCHES

Stephen Pankavich received his Ph.D. in Mathematical Sciences from Carnegie Mellon University in 2005 and has been interested in applications of mathematics throughout his career. He is currently an Assistant Professor at the Colorado School of Mines where he actively seeks to relate mathematics to the daily life of his undergraduate students while stressing its immense utility within science and engineering.

Rebecca Swanson received her Ph.D. in Mathematics from Indiana University in 2010. Prior to arriving at the Colorado School of Mines in 2012, she was an Assistant Professor at Nebraska Wesleyan University. She enjoys the beauty and elegance of discrete mathematics and the challenge of pedagogical innovation. In her free time, she also enjoys baking, running, and reading.