

Stripey Squares

Kelly Delp
 Ithaca College
 Ithaca, NY 14850, USA
 kelly.delp@gmail.com

Abstract

A black and white image can be represented by a matrix containing numbers between 0 and 1; we describe a way to visualize the squares of these image matrices, and more generally any non-negative power. The squares all share a common feature, stripes, and we explain the mathematics behind these.

Introduction

This paper grew out of explorations in a Linear Algebra course. Most exercises found in a typical Linear Algebra textbook require students to work with matrices of size $n \times m$, where each dimension is at most 6. Images provide good examples of objects that can be represented by large matrices. A black and white photo can be represented by an array of numbers between 0 (black) and 1 (white). One can ask, what does the sum of two images look like? Of course, when we add two image matrices together, the maximum value of the sum is likely greater than one. Averaging the entries gives a way to visualize the sum. Figure 1 contains an image which illustrates this (c), as well as some intermediate linear combinations (b and d).




Figure 1: *Linear Combinations*

After addition, the next operation studied by students is multiplication. A natural question arises.

Question 1. *What does the square of an image look like?*

We will give one potential answer to this question, and then look at more general powers of images.

Images of Squares

Suppose A is an $n \times n$ matrix corresponding to a black and white image. Since all of the entries in A are non-negative, A^2 will also have all non-negative entries. Also, since the value of any entry in A is at most 1, the maximum entry in A^2 is less than or equal to n . To create an image matrix we must rescale the values of A^2 so that they all fall between 0 and 1. Let M denote the maximum value of the entries in A^2 . Any increasing function $f : [0, M] \rightarrow [0, 1]$ could be applied to the entries of the matrix to create an image, with different functions creating different effects. For the purposes of this paper, we will use the linear rescaling

function, $f(x) = \frac{1}{M}x$. Simply put, whenever we have a matrix with non-negative values, some of which may be greater than one, we divide every entry in the matrix by the maximum entry in that matrix.

Experimenting with several images, we see that squares of image matrices all share similar characteristics—they are “stripey”. Figure 2 contains examples of squares of portraits. Despite being very different from the original images, some of the features, such as eyes, mouths and mustaches, seem to contribute to the square image.

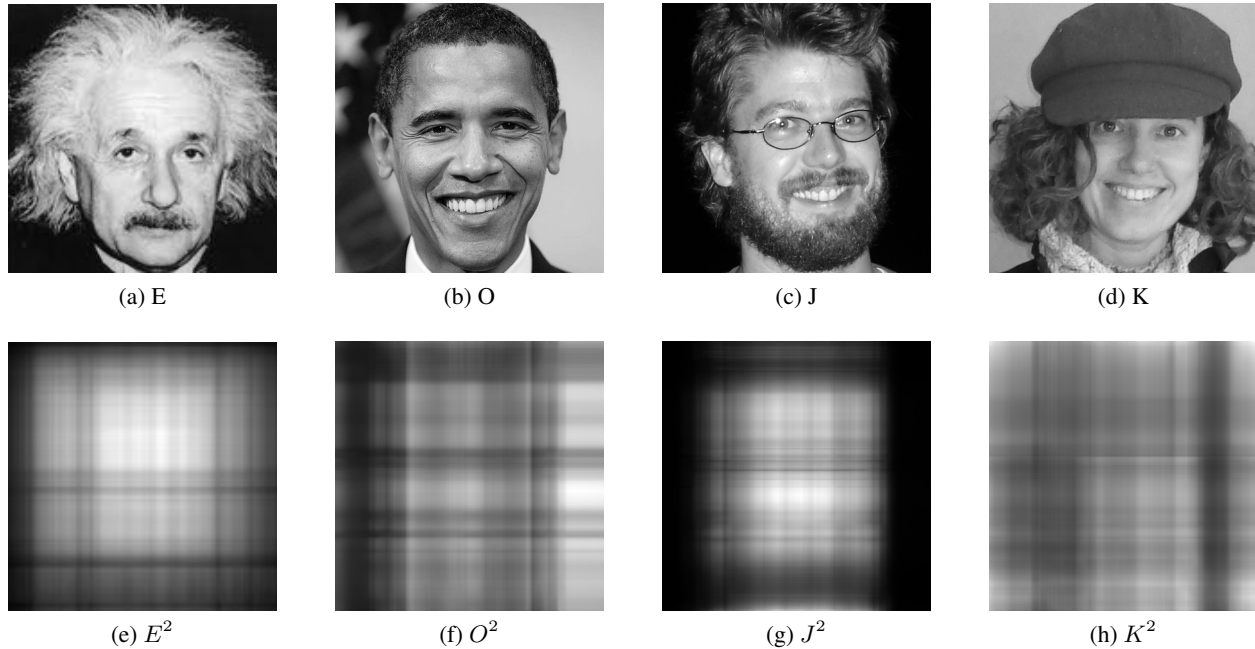


Figure 2: *Squares*

The image squares suggest the question: “Where do the stripes come from?” Although only portraits are shown above, we’ll argue that any image obtained by squaring will be striped. For evidence, consider the random (150×150) matrix and its square in Figure 3.

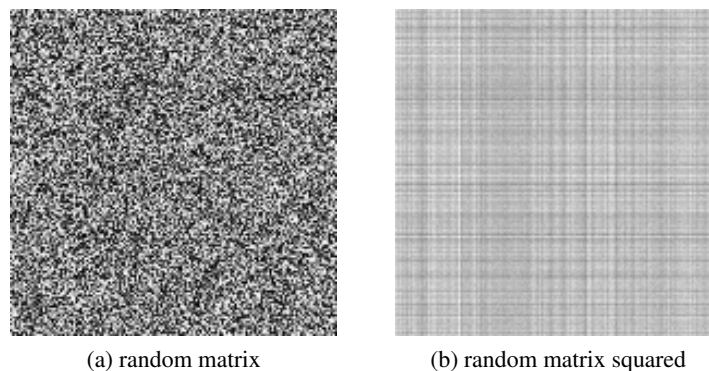


Figure 3: *Randomness and order*

The stripes in the squared image are a consequence of the Perron-Frobenius theorem, which is a statement about the eigenvalues and eigenvectors of positive matrices. Theorem 1 contains one version of the Perron-Frobenius theorem, which can be found in many linear algebra texts. See [1] for example. A *positive*

matrix is one in which all entries are real and positive. Matrices corresponding to the images can always be represented by positive matrices. Recall that a vector v is an eigenvector of A , with corresponding eigenvalue λ , if $Av = \lambda v$.

Theorem 1 (Perron-Frobenius). *Let A be a square positive matrix. Then A has a real eigenvalue λ_1 with the following properties:*

1. $\lambda_1 > 0$
2. λ_1 has a corresponding positive eigenvector, which is unique up to scale.
3. If λ is any other eigenvalue of A , then $|\lambda| < \lambda_1$.

We will sometimes refer to the Perron-Frobenius eigenvalue (eigenvector) as the *dominant* eigenvalue (eigenvector). Before discussing the stripe effect, we will consider a matrix of smaller dimension and discuss a consequence of Theorem 1.

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

The largest eigenvalue of this matrix is 5, and $v = (1, 1, 1)$ is the corresponding eigenvector. Every vector in the positive octant of \mathbb{R}^3 can be projected (by rescaling) into the larger triangle shown in Figure 4, which is defined by the intersection of the plane defined by $x + y + z = 1$ and the positive octant. The specific plane into which we are projecting is not important; this plane was chosen for visualization purposes. The important point is to think of vectors in this octant up to positive scale factors.

The point $p = (1/3, 1/3, 1/3)$ shown in Figure 4 corresponds to the eigenvector v . Since v is the dominant eigenvector, for any vector w in \mathbb{R}^3 with positive entries, applying A to w pulls w closer to the direction of v . In other words, after projecting Aw into the plane $x + y + z = 1$, the point Aw is closer to p . In Figure 4 we show two triangles; the larger triangle T is defined by the standard basis vectors $(1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$. The smaller triangle, $A(T)$, is the (projected) image of the larger triangle after applying the matrix A . The entire large triangle is mapped into the smaller one, which contains the eigenvector v .




Figure 4: Action of A on \mathbb{R}^3

Now let's return to an $(n \times n)$ image matrix A . For context, we note that all of the examples shown in this paper have resolution between $n = 150$ and $n = 625$. Since all of the entries in the matrix A are positive, the Perron-Frobenius theorem tells us the largest eigenvalue is a positive real number, and the corresponding eigenvector v has all positive entries as well. Suppose $A = [a_1|a_2|a_3|\dots|a_n]$, where a_i denotes the i^{th}

column vector of A . We may think of A^2 as

$$A^2 = AA = A [a_1|a_2|a_3|\cdots|a_n] = [Aa_1|Aa_2|Aa_3|\cdots|Aa_n].$$

As in the three dimensional case, for any of the columns a_i , the vector Aa_i lies closer to the direction of the eigenvector v . So, after rescaling by the maximum value of A^2 , the column vector of A^2 which corresponds to Aa_i has values which are very close to the rescaled eigenvector v . Therefore, all of the column vectors of the rescaled A^2 are very similar, creating the horizontal stripes in the square image.

The vertical stripes can similarly be explained. Let \hat{a}_i denote a row of A , then i^{th} row of A^2 is given by $\hat{a}_i A$. Let w be the *left* eigenvector of A corresponding to the largest eigenvalue. (Note that the left eigenvectors of a matrix are equal to the eigenvectors of the transpose of A .) After rescaling, all the rows of A^2 are close to the rescaled vector w , creating the vertical stripes.




(a) F (b) Matrix with the dominant eigenvector of F^T in every row(c) Matrix with the dominant eigenvector of F in every column(d) F^2

Figure 5: *Matrices of an eigenvector*

The above discussion is best illustrated with pictures. Figure 5a contains an image of this author appreciating a work by artist Parastou Forouhar, taken in the Queensland Gallery of Modern Art in Brisbane, Australia. Let F denote the corresponding image matrix. Figure 5d contains an image created from F^2 . In 5c we show the image corresponding to the matrix that has the rescaled Perron-Frobenius eigenvector of F in every column. Compare to F^2 and we see the eigenvector coincides with the horizontal stripes in the square matrix. The image in 5b is from the matrix which has the (rescaled) Perron-Frobenius eigenvector of F^T in every row, which clearly corresponds to the vertical stripes.

A natural question arises: What does the cube (or higher power) of an image matrix look like? The answer is that not very much interesting happens after the square. Essentially, the sequence of rescaled power matrices converges to the rescaled product of the matrices found in Figure 5c and 5b, each of which contains a single eigenvector repeated. The convergence happens very quickly. This is evident from how closely the eigenvector matrices agree with the square image illustrated in Figure 5. It's almost impossible to detect with the naked eye the difference in even the "square" image and the limiting product matrix. All of the higher powers also produce images that look like the image found in 5d.

Visualizing subgroups of $GL(n, \mathbb{C})$

For any integer n , the set of $n \times n$ matrices with entries in a field F and non-zero determinant forms a group called the general linear group, $GL(n, F)$. If $F = \mathbb{R}$ this group is also a manifold of dimension n^2 . In the introduction we described a way to interpolate between two matrices; this interpolation (Figure 1) can be thought of as a straight line in the group $GL(150, \mathbb{R})$, with start and end points giving the matrices for the original images. In this section we'll describe a way to visualize another sort of path in these large dimensional (Lie) groups.

Let M be a diagonalizable matrix, so $M = V^{-1}DV$ where V is a matrix containing the eigenvectors of M , and D is the matrix containing the eigenvalues of M on the diagonal. Using the diagonalization we may define non-integer powers of M by $M^t = V^{-1}D^tV$, where D^t is a diagonal matrix containing the t^{th} power of the eigenvalues. For diagonalizable matrices, M^t defines a one-parameter subgroup of $GL(n, F)$. We will describe a way to visualize a path in this subgroup when M is a matrix corresponding to an image, the results of which are shown in Figure 6.

The matrix M that we begin with comes from the image found in Figure 6i. Although this matrix is not diagonalizable over \mathbb{R} , it is diagonalizable over the complex numbers. (Thus far every image matrix we have checked has been diagonalizable over \mathbb{C} .) Therefore for non-integer values of t , the matrix M^t has complex numbers. To create the images found in Figure 6, we take the norm of all the entries in M^t , and then use the same normalization (divide by the max value of the matrix) to create an image. For $t = 0$ we get the identity matrix, and $t = 2$ is the previously discussed square.

This is just one way to visualize the matrices with complex values. Another method we have tried with reasonable success uses color; for each complex entry in the matrix M^t we assign the real part of an entry with a value on the red spectrum, and the complex part a green value. It's interesting to see where the numbers with the largest imaginary values appear and disappear along the paths $0 \leq t \leq 1$ and $1 \leq t \leq 2$.

Methods and Conclusions

All of the images in this paper were created in *Mathematica*, which has built-in support for image processing. For example, it is possible to "drag and drop" jpegs into a *Mathematica* notebook and immediately start manipulating. The ease with which one can turn matrices into images makes us hopeful that these tools can be used by undergraduates (with little to no programming knowledge) for class projects and exploration.

The mathematics used to explain the striped effect that occurs in the image can be found in a typical undergraduate linear algebra class. However, this author thinks it is unlikely that she would have predicted the effect. The images illuminated the mathematics, giving a nice illustration of the Perron-Frobenius Theorem.

References

- [1] David Poole. Linear Algebra: A Modern Introduction, 3rd ed. *Cengage Learning*, (2010).

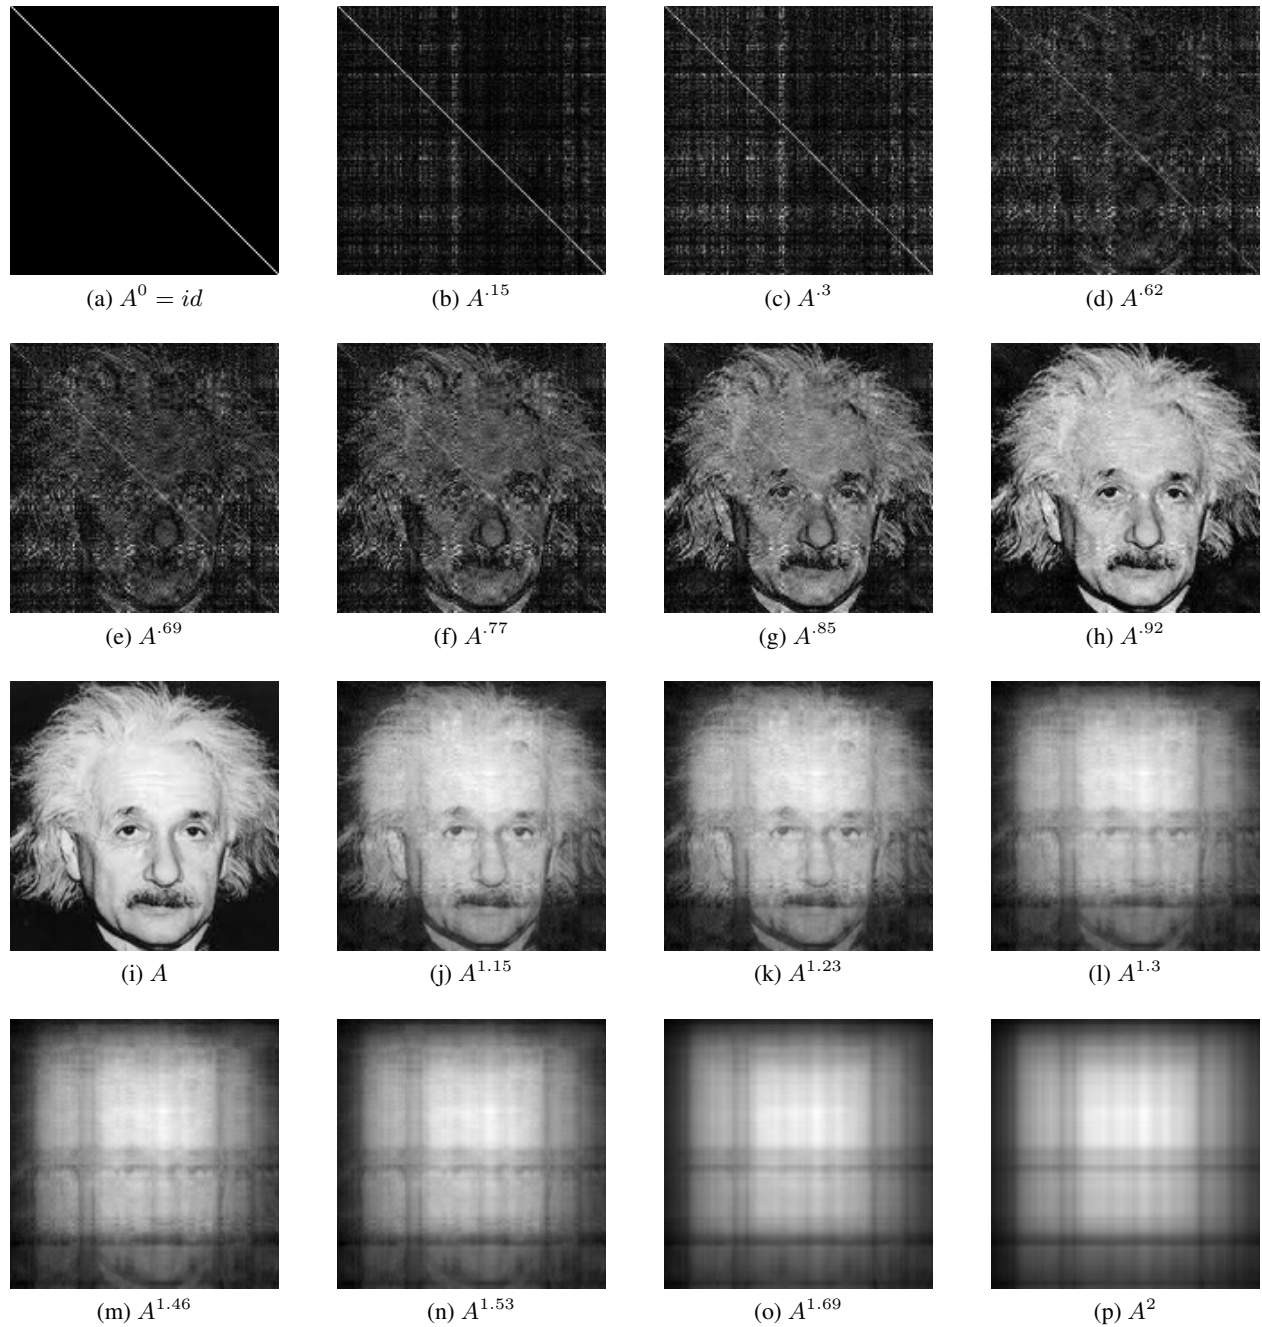


Figure 6: A^t for $0 \leq t \leq 2$