

Essential Math for Data Science: Matrix Diagonalization Clearly Explained

Irene Markelic

December 18, 2025




Figure 1: Image generated using AI (Gemini)

1 Introduction

Matrix diagonalization is a crucial concept in linear algebra with tons of applications. Even more importantly, it is the KEY to understanding advanced and essential data science topics like the singular value decomposition (SVD) and principal component analysis (PCA).

This article goes beyond intuition building only - it provides you with a very accessible mathematical derivation of the topic, making it perfectly suitable for beginners while also

erving as a knowledge refresher for seasoned professionals. It is meant as your friendly math guide.

So what is Matrix diagonalization all about? It allows us to rewrite a matrix \mathbf{A} in a simpler form involving a diagonal matrix \mathbf{D} . This simplified form makes many otherwise difficult computations - like raising a matrix to a high power - trivial. Understanding matrix diagonalization is not difficult per se, but it builds on various other concepts. The following articles, written as supplementary material, provide a detailed treatment of these concepts.

- matrix multiplication
- linear transformations vs change of basis and
- eigenvalues and eigenvectors.

2 Deriving the Diagonalization Formula: $\mathbf{AX} = \mathbf{X}\Lambda$

In this section we will derive the fundamental matrix expression $\mathbf{AX} = \mathbf{X}\Lambda$ which is the cornerstone of diagonalization.

We start with the definition of an eigenvalue-eigenvector pair (\mathbf{x}, λ) for a square matrix \mathbf{A} :

$$\mathbf{Ax} = \lambda\mathbf{x} \tag{1}$$

This equation means that applying the transformation \mathbf{A} to the vector \mathbf{x} simply results in scaling \mathbf{x} by the scalar λ , without changing its direction.

For an $n \times n$ matrix \mathbf{A} we can potentially compute all the n eigenvalue-eigenvector pairs. We can then take all the eigenvectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, and stack them into a matrix, which we'll call \mathbf{X} (capital letter, to denote that its a matrix of vectors). For example, for

$$\mathbf{A} = \begin{bmatrix} 4 & 1 \\ 2 & 3 \end{bmatrix} \tag{2}$$

we can compute the two eigenvalue-eigenvector pairs $\lambda_1 = 5$, $\mathbf{x}_1 = [0.7, 0.7]^T$ and $\lambda_2 = 2$, $\mathbf{x}_2 = [-0.4, 0.8]^T$. We can then form the matrix \mathbf{X} by using these eigenvectors as its columns and we choose to start with the eigenvector that has the largest eigenvalue, \mathbf{x}_1 and then we add the second eigenvector with the second-largest eigenvalue, \mathbf{x}_2 .

$$\mathbf{X} = [\mathbf{x}_1 \quad \mathbf{x}_2] = \begin{bmatrix} 0.7 & -0.4 \\ 0.7 & 0.8 \end{bmatrix} \tag{3}$$

Next, we create a diagonal matrix, $\mathbf{\Lambda}$ (capital Lambda), which contains the corresponding eigenvalues on its main diagonal, in the **exact same order as their eigenvectors in \mathbf{X}** :

$$\mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} = \begin{bmatrix} 5 & 0 \\ 0 & 2 \end{bmatrix}$$

While sorting the eigenvalues from largest to smallest, as we did here, is crucial in certain applications (like Principal Component Analysis), the fundamental mathematical requirement is simply consistent ordering. That is, if you change the order of the columns in \mathbf{X} (e.g., swapping \mathbf{x}_1 and \mathbf{x}_2), you must also swap the corresponding diagonal entries in $\mathbf{\Lambda}$ (i.e., λ_1 and λ_2) to maintain this correspondence.

With these definitions, the following fundamental expression is true — and we will show why:

$$\mathbf{AX} = \mathbf{X}\mathbf{\Lambda} \tag{4}$$

Note the order: it is $\mathbf{X}\mathbf{\Lambda}$, not the other way round.

To prove this identity, we will look at the result of the matrix multiplication on both sides of the equation and compare the columns.

2.1 The Left Side: \mathbf{AX}

For the left side recall a key property of matrix multiplication: when a matrix \mathbf{X} is multiplied by a matrix \mathbf{A} , the resulting matrix's columns are equal to \mathbf{A} times the corresponding column of \mathbf{X} (see here, for a refresher on matrix multiplication).

$$\mathbf{AX} = [\mathbf{Ax}_1 \quad \mathbf{Ax}_2 \quad \dots \quad \mathbf{Ax}_n] \tag{5}$$

Now, we apply the definition of the eigenvalue-eigenvector pair, $\mathbf{Ax} = \lambda\mathbf{x}$, to each column:

$$\mathbf{AX} = [\lambda_1\mathbf{x}_1 \quad \lambda_2\mathbf{x}_2 \quad \dots \quad \lambda_2\mathbf{x}_n] \tag{6}$$

2.2 The Right Side: $\mathbf{X}\mathbf{\Lambda}$

For the right side, we use another fundamental property of matrix multiplication: multiplying a matrix \mathbf{X} by a diagonal matrix $\mathbf{\Lambda}$ scales the columns of \mathbf{X} by the corresponding diagonal elements of $\mathbf{\Lambda}$. To see this, let's look at the first column of the product $\mathbf{X}\mathbf{\Lambda}$ for our $n \times n$ case:

The first column of $\mathbf{\Lambda}$ is the vector $[\lambda_1, 0, \dots, 0]^T$. The multiplication \mathbf{X} times this column

gives:

$$\mathbf{X} \begin{bmatrix} \lambda_1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \lambda_1 \mathbf{x}_1 + 0\mathbf{x}_2 + \dots + 0\mathbf{x}_n = \lambda_1 \mathbf{x}_1 \quad (7)$$

This process is repeated for every column of $\mathbf{\Lambda}$. The i -th column of $\mathbf{\Lambda}$ is a vector with λ_i in the i -th position and zeros elsewhere. Therefore, the i -th column of $\mathbf{X}\mathbf{\Lambda}$ simplifies to $\lambda_i \mathbf{x}_i$. Stacking these resulting columns gives us:

$$\mathbf{X}\mathbf{\Lambda} = [\lambda_1 \mathbf{x}_1 \quad \lambda_2 \mathbf{x}_2 \quad \dots \quad \lambda_n \mathbf{x}_n] \quad (8)$$

By comparing equation (5) and equation (7), we see that the columns of $\mathbf{A}\mathbf{X}$ are identical to the columns of $\mathbf{X}\mathbf{\Lambda}$. Thus, it is clear that $\mathbf{A}\mathbf{X} = \mathbf{X}\mathbf{\Lambda}$ is true.

3 Deriving the Diagonalization Formula: $\mathbf{A} = \mathbf{X}\mathbf{\Lambda}\mathbf{X}^{-1}$

We established the fundamental matrix relationship $\mathbf{A}\mathbf{X} = \mathbf{X}\mathbf{\Lambda}$ in the previous section. This expression is crucial for understanding matrix transformations and is the starting point for diagonalization. Now, let's manipulate it further to isolate \mathbf{A} and $\mathbf{\Lambda}$. We start with this expression:

$$\mathbf{A}\mathbf{X} = \mathbf{X}\mathbf{\Lambda}$$

If, and only if, the matrix \mathbf{X} is invertible (meaning \mathbf{X}^{-1} exists), we can post-multiply both sides by \mathbf{X}^{-1} to isolate \mathbf{A} :

$$\mathbf{A} = \mathbf{X}\mathbf{\Lambda}\mathbf{X}^{-1} \quad (9)$$

Alternatively, we can pre-multiply both sides of the original equation by \mathbf{X}^{-1} to isolate the diagonal matrix $\mathbf{\Lambda}$:

$$\mathbf{\Lambda} = \mathbf{X}^{-1}\mathbf{A}\mathbf{X} \quad (10)$$

3.1 When Can We Diagonalize? The Invertibility of \mathbf{X}

This manipulation immediately raises a critical question: Can we always compute \mathbf{X}^{-1} ? The answer is no. The inverse \mathbf{X}^{-1} can only be computed when \mathbf{X} is invertible. A matrix is invertible if and only if all of its columns (in this case, the eigenvectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$) are mutually linearly independent. A square $n \times n$ matrix is guaranteed to have n linearly independent eigenvectors if and only if for every eigenvalue, the geometric multiplicity equals the algebraic multiplicity, [1] p. 77. (The algebraic multiplicity is the number of times an eigenvalue is a repeated root of the characteristic polynomial (e.g., $\lambda_i = \lambda_j = 2$).

The geometric multiplicity is the number of linearly independent eigenvectors associated with that eigenvalue. We always have Geometric Multiplicity \leq Algebraic Multiplicity [2] p. 311.) A sufficient condition that guarantees this is if all of its n eigenvalues have distinct values (see here for more information or better a standard textbook on linear algebra). Therefore, the diagonalization formulae hold only for those special square matrices that possess n linearly independent eigenvectors.

3.2 A Geometrical Interpretation of Diagonalization

We have seen that if the conditions for \mathbf{X}^{-1} are met, we can rewrite the original matrix \mathbf{A} as the product of three other matrices:

$$\mathbf{A} = \mathbf{X} \cdot \mathbf{\Lambda} \cdot \mathbf{X}^{-1}$$

Since we have derived this formula we do haven an abstract understanding of it, however, there is a very simple geometric interpretation, too and understanding this aspect in addition to the mathematical derivation will deepen our understanding of this process. Prerequisite is to have a solid understanding of linear transformations and change of basis as they work beautifully together in the process of **diagonalization**:

So here we go: We know that:

- \mathbf{D} is a diagonal matrix. Diagonal matrices are highly desirable because they lead to greatly simplified computations (e.g., raising a matrix to a power).
- \mathbf{P} is an invertible matrix whose columns are \mathbf{A} 's eigenvectors. This matrix is also known as \mathbf{A} 's **eigenbasis**.
- \mathbf{P}^{-1} is the inverse of the matrix of eigenvectors \mathbf{P} .

Let's say we apply \mathbf{A} to some vector \mathbf{x} . With diagonalization we can write this as:

$$\mathbf{Ax} = \mathbf{XDX}^{-1}\mathbf{x} \tag{11}$$

And using the associativity of matrix multiplication we can compute this in the following order: $(\mathbf{X}(\mathbf{D}(\mathbf{X}^{-1}\mathbf{x})))$ We can see that this way diagonalization breaks the transformation down into three linked, simpler operations:

1. $\mathbf{X}^{-1}\mathbf{x}$: This is a change of basis. This means we transform the input vector \mathbf{x} from the standard basis coordinates into the coordinates of \mathbf{A} 's eigenbasis.

$$\mathbf{x}_{\text{eigenbasis}} = \mathbf{X}^{-1}\mathbf{x}$$

2. $\mathbf{D}\mathbf{x}_{\text{eigenbasis}}$: This is a simple transformation. We simply apply the diagonal matrix (\mathbf{D}) to the new coordinates. Applying a diagonal matrix only involves a scaling of the components, making it computationally cheap.

$$\mathbf{y}_{\text{eigenbasis}} = \mathbf{D}(\mathbf{X}^{-1}\mathbf{x})$$

3. $\mathbf{Xy}_{\text{eigenbasis}}$: This is another change of basis. We use the original basis matrix (\mathbf{X}) to transform the result back from the eigenbasis coordinates into our standard, familiar coordinate system.

$$\mathbf{y}_{\text{standard}} = \mathbf{X}(\mathbf{D}(\mathbf{X}^{-1}\mathbf{x}))$$

In summary: Diagonalization can be interpreted as first transforming the input (\mathbf{x}) into the matrix's (\mathbf{A}) eigenbasis. In this special representation the transformation reduces to a simple stretch of the input-vector into the various dimensions defined by the eigenvectors of \mathbf{A} . And then we simply transform the result back into the standard basis.

4 Summary of Keypoints

Congratulations! Let's review the crucial points that we've covered:

- We have derived the formula for matrix diagonalization to be $\mathbf{Ax} = \mathbf{XDX}^{-1}\mathbf{x}$.
- We know that at its core is \mathbf{A} 's eigenbasis. If this can't be computed, diagonalization with this method is not possible. You know that diagonalization is only possible for $n \times n$ square matrices where the following condition holds for every eigenvalue λ : Geometric Multiplicity of $\lambda = \text{Algebraic Multiplicity of } \lambda$. The simplest sufficient condition that guarantees this is for the matrix to have n distinct eigenvalues.
- Because of the above we understand the importance of eigenvalues and eigenvectors for matrix diagonalization.
- We understand that diagonalization has a geometric interpretation that involves expressing the input in terms of \mathbf{A} 's eigenvectors, doing a then-simple stretch, and a transformation back into the standard basis.

While diagonalization is an excellent tool, **its utility is limited to square matrices that have a complete and distinct set of eigenvalues.**

So, where to go from here? Having established these fundamental concepts, the next major topic is the singular value decomposition (SVD). SVD represents a major leap forward in linear algebra. Essentially, SVD is a powerful generalization of diagonalization, enabling us to decompose any matrix—not just square ones. It is a more advanced, more general, and ultimately more versatile technique.

References

- [1] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. 2nd. Cambridge; New York: Cambridge University Press, 2013. ISBN: 9780521839402.
- [2] Gilbert Strang. *Introduction to Linear Algebra, Fifth Edition*. Wellesley-Cambridge Press, 2016.