Linear Algebra I: Matrix Representations of Linear Transforms

Nathan Ratliff

Oct 17, 2014

Abstract

Vector spaces exist independent of our choice of coordinate representation. Your house, the arrangement of its furniture, the cat sitting on the windowsill, or the glass falling from your hand toward its impending explosion across the floor, all exist independent of whether you choose to describe it all relative to a coordinate system aligned with your room or one aligned align with napping neighbor John at rest on his bed. This document introduces the notion of an abstract vector space and how coordinates, and ultimately matrices of linear transformations, arise from the underlying linearity structures imposed on the system.

1 Vector spaces

This document assumes basic familiarity with linear algebra from a coordinate centric perspective. Many introductory texts present linear algebra as the algebra of matrices that transform vectors in $\mathbb{R}^n$ to vectors in $\mathbb{R}^m$. That's an excellent way to gain intuition for the problem, but when we start talking about physical systems that exist in the real-world, there's something arbitrary about a specific choice of coordinates. The space around us exists independent of how we choose to describe it, and we should be careful that algorithms we derive and run aren't reliant of the specifics of something as arbitrary as that choice of coordinates. My algorithm for getting a robot to navigate around your house shouldn't have anything to do with the fact that I've decided to align the z-axis upward and the x and y axes with latitude and longitude. I probably chose those coordinates just because it's easier for me to do the calculations or because it's nice for transforming GPS information into something the robot can use. But it'd be a shame if that choice ended up defining the behavior of the robot.[1]

It's conceptually useful, and in many cases, practically important, to define what a vector space is independent of a coordinate system. The coordinate system can be chosen over the top of that space in the end after we've derived the algorithm and when need to actually do some real computations. But every

---

[1] There certainly are algorithms that do have such a dependence, but we try to avoid that as much as possible.

time we choose a coordinate system, we should remain aware of the underlying vector space and its geometry.

This section introduces the basic definitions and properties of abstract vector spaces and what we mean by "coordinates" in the space. Section 2 then shows how matrices and matrix multiplication arise as what can be viewed as a bookkeeping technique for tracking how coordinates in one space transform to coordinates in another space under a linear transformation.

We introduce the notation $|a\rangle$ to denote an abstract vector. This notation was introduced by Paul Dirac in 1939 Hassani (2013); Wikipedia (2002-2014) and is used quite a bit in physic where the study of physical phenomena dictates coordinate independent representations. We'll follow that convention for much of our discuss of coordinate-free linear algebra to emphasize the difference between an abstract vector and a collection of coordinates representing the vector. A good reference for the concepts presented here is Sadri Hassani's Mathematical Physics text Hassani (2013).

1.1   Basic definition

We say a collection of elements $\mathcal{V}$ is a vector space over the real numbers $\mathbb{R}$ if it satisfies the following properties

1. **Summation.** There is an operation $+$, called *summation*, such that for every pair of vectors $|a\rangle$ and $|b\rangle$ in $\mathcal{V}$, there is another vector $|a\rangle + |b\rangle$ also in $\mathcal{V}$. The summation operation must have the following properties:

   (a) *Commutativity.* $|a\rangle + |b\rangle = |b\rangle + |a\rangle$.

   (b) *Associativity.* $|a\rangle + (|b\rangle + |c\rangle) = (|a\rangle + |b\rangle) + |c\rangle$.

   (c) *Zero vector.* There exists a vector $|0\rangle \in \mathcal{V}$ such that $|a\rangle + |0\rangle = |a\rangle$ for all $|a\rangle \in \mathcal{V}$.

   (d) *Additive inverse.* For every $|a\rangle \in \mathcal{V}$, there is a corresponding element $|-a\rangle$ such that $|a\rangle + |-a\rangle = |0\rangle$.

2. **Scaling.** For every $\alpha \in \mathbb{R}$ and $|a\rangle \in \mathcal{V}$, there is a corresponding vector $\alpha|a\rangle \in \mathcal{V}$ such that

   (a) Order of scaling doesn't matter: $\alpha(\beta|a\rangle) = (\alpha\beta)|a\rangle$.

   (b) Scaling by 1 does nothing: $1|a\rangle = |a\rangle$.

3. **Distributive property.** Multiplication of scalars and vectors is distributive.

   (a) Across vectors: $\alpha(|a\rangle + |b\rangle) = \alpha|a\rangle + \alpha|b\rangle$.

   (b) Across scalars: $(\alpha + \beta)|a\rangle = \alpha|a\rangle + \beta|a\rangle$.

This definition seems like a lot of details, but the bottom line is that all of the interesting structural properties of vector spaces pertinent to linear algebra rely only on this list of requirements. But more broadly, these requirements

make sense. For a vector space, we need to be able to add elements together in a nice way (the properties of summation show that the vectors form a *group* under addition), and we need to be able to scale the vectors by real values (the properties of scaling define what we mean when we say scaling behaves intuitively—the value of 1 retains the role of the *field's* multiplicative identity and the order we multiply doesn't matter). Moreover, addition and multiplication between scalars and vectors behaves the same way it would under scalars alone (distributive property). These terms *group* and *field* are interesting for those with experience with Abstract Algebra Dummit & Foote (2003); Hassani (2013), but for the most part here, this will be the last we mention of it. The main point of this discussion is that there are some theoretically sound reasons why this collection of properties arise.

1.2   Linear combinations of vectors

In short, a vector space is a set of objects that can be added together and scaled in nice ways. These "niceness" properties, such as the existence of a zero vector and the distributive property, form the core structure in linear algebra that we're able to exploit. Much of the interesting structure centers around basis representations; we state many of the key results here without proof.

Since we can add scaled vectors together, we can talk about representing a vector as a **linear combination** of other vectors. Suppose we have a collection of vectors $\mathcal{B} = \{|\boldsymbol{b}_1\rangle, \ldots, |\boldsymbol{b}_n\rangle\}$. Then a vector that can be expressed as

$$|\boldsymbol{v}\rangle = \sum_{i=1}^{n} \alpha_i |\boldsymbol{b}_i\rangle \tag{1}$$

is said to be a linear combination of the vectors in $\mathcal{B}$. The set of all vectors that can be expressed this way is the **span** of $\mathcal{B}$:

$$\mathbf{span}\{\mathcal{B}\} = \left\{ |\boldsymbol{v}\rangle \;\middle|\; |\boldsymbol{v}\rangle = \sum_{i=1}^{n} \alpha_i |\boldsymbol{b}_i\rangle \text{ for some } \alpha_i \right\}. \tag{2}$$

It's straightforward to prove that the span of a set of vectors is itself a vector space. Note that $\mathcal{B}$ can be highly redundant. Any vector that's already in the span of $\mathcal{B}$ can, itself, be added to $\mathcal{B}$ and the space spanned by the set doesn't change. It's just redundant.

So perhaps a more interesting set is one that isn't redundant. A minimal set of vectors $\mathcal{B}$ that spans a vector space $\mathcal{V}$ is said to be a **basis** for that space. This minimality property says that no vector $|\boldsymbol{b}_i\rangle \in \mathcal{B}$ can be expressed as a linear combination of other vectors in $\mathcal{B}$ (otherwise, $\mathcal{B}$ wouldn't be minimal). When $\mathcal{B}$ is minimal (i.e. it's a **basis** for its span), we say that its elements are **linearly independent**.

Many text books define linear independence as the property

$$\sum_{i=1}^{n} \alpha_i |\boldsymbol{b}_i\rangle = |\boldsymbol{0}\rangle \quad \Rightarrow \quad \alpha_i = 0 \ \text{ for all } i. \tag{3}$$
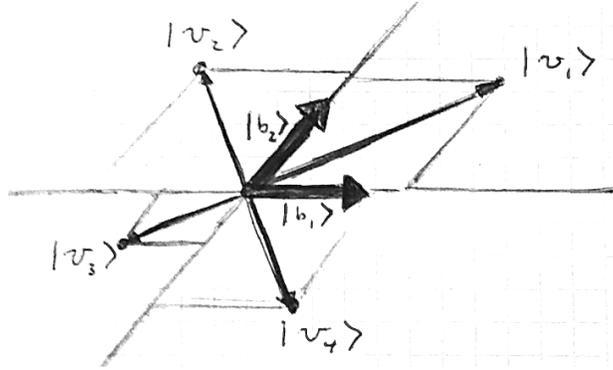
Figure 1:   Coordinate representation of 2-dimensional vectors in the span of a basis $\{|\boldsymbol{b}_1\rangle, |\boldsymbol{b}_2\rangle\}$.

Our descriptive definition above (no element is a linear combination of the others) is equivalent and perhaps more intuitive. You can think of it either way.

Finally, now that we have a definition of a minimal set $\mathcal{B}$ spanning a vector space $\mathcal{V}$ (i.e. a basis for $\mathcal{V}$), we now know something about the size of the space. The minimal number of vectors necessary to represent an entire vector space is called the **dimension** of that space. In other words, if $\mathcal{B}$ is a basis for $\mathcal{V}$, and $|\mathcal{B}| = n$, we say that $\mathcal{V}$ is n-dimensional. Importantly, it can be shown that the dimensionality of a space is unique. All bases spanning $\mathcal{V}$ are going to have n elements. Figure 1 depicts the vectors lying in the span of a basis for a 2-dimensional space.

1.3   Coordinate representations of vectors in a basis

One concept of importance to this document is the notion of a **coordinate representation** of vectors in a vector space. Traditionally, we view some of the canonical vector spaces, such as $\mathbb{R}^n$, as a collection of coordinates. We have a series of axes (the x-axis, y-axis, and z-axis for $\mathbb{R}^3$), and we can represent all points via their coordinates along these axes. So we say that the point $(1.2,\ 3,\ -1)$ is in $\mathbb{R}^3$, for instance, but that is to say that the point which can be represented as 1.2 along the unit vector $\hat{\boldsymbol{x}}$, 3 along the unit vector $\hat{\boldsymbol{y}}$, and -1 along the unit vector $\hat{\boldsymbol{z}}$, is in $\mathbb{R}^3$. Coordinates specify the unique representation of a point in a vector space with respect to a given basis.

In the general case, when we use a basis $\mathcal{B}$ to represent an abstract vector space $\mathcal{V}$, we say that the coefficients $(\alpha_1, \ldots, \alpha_n)$ of the unique expansion of a vector $|\boldsymbol{v}\rangle = \sum_i \alpha_i |\boldsymbol{b}_i\rangle$ in $\mathcal{V}$ form the coordinate representation of $|\boldsymbol{v}\rangle$ with respect to the basis $\mathcal{B}$.

2   Linear transforms and matrix representations

All of the above definitions are used to just describe the structure of a vector space and give us tools to represent vectors as coordinates with respect to some basis. But as of yet there's no real algebra. Just structure. The algebra appears when we start to discuss transformations from one space to another. These ideas arise as a generalization of matrix algebra, where we study how to exploit the geometry of systems of linear equations in $\mathbb{R}^n$ to learn something about the space of solutions. We'll see here that in this abstract setting we can think of a matrix as a representation of a linear transformation given a choice of bases for the domain and range spaces.[2] The matrix representation tells us how the coordinates of a vector with respect to the domain's basis transform to the coordinates of the transformed vector with respect to the range's basis.

To formalize these ideas, let $\mathcal{V}$ and $\mathcal{U}$ be vector spaces of dimension n and m, respectively. And suppose $\boldsymbol{T} : \mathcal{V} \to \mathcal{U}$ is a transform mapping vectors $|\boldsymbol{v}\rangle \in \mathcal{V}$ to a vector $\boldsymbol{T}|\boldsymbol{v}\rangle \in \mathcal{U}$. Then we say that $\boldsymbol{T}$ is a **linear transform** if transforming a linear combination of vectors is the same as linearly combining the transformed vectors. Combining then transforming is the same as transforming then combining. Formally, for any $\alpha, \beta \in \mathbb{R}$ and $|\boldsymbol{a}\rangle, |\boldsymbol{b}\rangle \in \mathcal{V}$, the following relation holds

$$\boldsymbol{T} \Big(\alpha|\boldsymbol{a}\rangle + \beta|\boldsymbol{b}\rangle \Big) = \alpha\boldsymbol{T}|\boldsymbol{a}\rangle + \beta\boldsymbol{T}|\boldsymbol{b}\rangle. \tag{4}$$

Lets take a look at what happens to the coefficients of basis representations under a linear transform. First, we need to choose bases for the domain and range. Let $\mathcal{B}_{\mathcal{V}} = \{\boldsymbol{b}_j\}_{j=1}^n$ be a basis for the domain space $\mathcal{V}$ and let $\mathcal{B}_{\mathcal{U}} = \{\boldsymbol{a}_i\}_{i=1}^m$ be a basis for the range space $\mathcal{U}$.

An arbitrary vector $|\boldsymbol{v}\rangle \in \mathcal{V}$ is represented as a linear combination of basis elements in the domain. Specifically, $|\boldsymbol{v}\rangle = \sum_{j=1}^n \alpha_j|\boldsymbol{b}_j\rangle$ where $\alpha_j \in \mathbb{R}$. Since $\boldsymbol{T}$ is linear, we can write the transformed vector as

$$\boldsymbol{T}|\boldsymbol{v}\rangle = \sum_{j=1}^n \alpha_j \boldsymbol{T}|\boldsymbol{b}_j\rangle. \tag{5}$$

Thus, if we just knew how to express each transformed $\boldsymbol{T}|\boldsymbol{b}_j\rangle$ as a linear combination of basis elements in the range space $\mathcal{U}$, we can get an expression for the coefficients of the transformed vector $\boldsymbol{T}|\boldsymbol{v}\rangle$ in range basis $\mathcal{B}_{\mathcal{U}}$ in terms of the coefficients of the original vector $|\boldsymbol{v}\rangle$ in domain basis $\mathcal{B}_{\mathcal{V}}$.

More concretely, let $(\beta_{1j}, \ldots, \beta_{mj})$ be the coordinates of $\boldsymbol{T}|\boldsymbol{b}_j\rangle$ in terms of $\mathcal{B}_{\mathcal{U}}$. Then $\boldsymbol{T}|\boldsymbol{b}_j\rangle = \sum_{i=1}^m \beta_{ij}|\boldsymbol{a}_i\rangle$ for each j. Plugging this expansion into

---

[2]Remember, the *domain space* is the space of inputs and the *range space* is the space of outputs. For a transform $\boldsymbol{T} : \mathcal{V} \to \mathcal{U}$, $\mathcal{V}$ is the domain and $\mathcal{U}$ is the range.

Equation 5 gives

$$\boldsymbol{T}|\boldsymbol{v}\rangle = \sum_{j=1}^{n} \alpha_j \sum_{i=1}^{m} \beta_{ij}|\boldsymbol{a}_i\rangle = \sum_{i=1}^{m} \underbrace{\left(\sum_{j=1}^{n} \beta_{ij}\alpha_j\right)}_{\gamma_i} |\boldsymbol{a}_i\rangle. \qquad (6)$$

The expression for the transformed coefficients γ_i here is very regular. Collecting all of these coefficients (γ_i, β_{ij}, and α_j) up into vector and matrix form gives the following relation:

$$\begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \vdots \\ \gamma_m \end{bmatrix} = \begin{bmatrix} \beta_{11} & \beta_{12} & \cdots & \beta_{1n} \\ \beta_{21} & \beta_{22} & \cdots & \beta_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \beta_{m1} & \beta_{m2} & \cdots & \beta_{mn} \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix}. \qquad (7)$$

In other words, matrix multiplication is the bookkeeping operation we use to track how coefficients in a domain expansion transform to coefficients in the range expansion. Importantly, for any linear transformation, there is an infinite number of matrix representations of the transformation. Each distinct pair of possible bases gives a different matrix.

Note that the coordinates of each domain basis basis element $|\boldsymbol{b}_j\rangle \in \mathcal{B}_\mathcal{V}$, represented in that same basis $\mathcal{B}_\mathcal{V}$, are given by the Kronecker delta: $\boldsymbol{\delta}_j$, which is zero for all coordinates i, except when $i = j$, at which point it is 1. (That's a fancy way of saying simply that $|\boldsymbol{b}_j\rangle = \sum_i \alpha_i|\boldsymbol{b}_i\rangle$ holds only when all the $\alpha'_k s$ are zero except for α_j, which must be 1.) Multiplying $\boldsymbol{\delta}_j$ by a matrix just extracts the jth column of that matrix. So in this case, the jth column must be the coordinate representation of the jth transformed basis element, as represented in the range basis.

3 An example: The matrix of a rotation in 2D

Suppose we have a 2D space. We typically think of the space in terms of Cartesian coordinates (x, y), but that's only a convention instilled in us since grade school. Really, it's just a linear space that could be described by any basis. To be concrete, we could even give it a physical manifestation: it could be the wall right next to you. So given a point on the wall, it's perfectly reasonable to talk about rotations around that point by an angle θ independent of any coordinate system (see Figure 2, left). In particular, we really don't want the definition of rotation to change based on the coordinate system we choose—the wall's geometry is independent of our choice of coordinate basis and the specific choice shouldn't be critical to the problem definition.

Of course, when we're computing transformations in 2D, we need create a tangible representation of the 2D space, so we need to choose a set of coordinates.
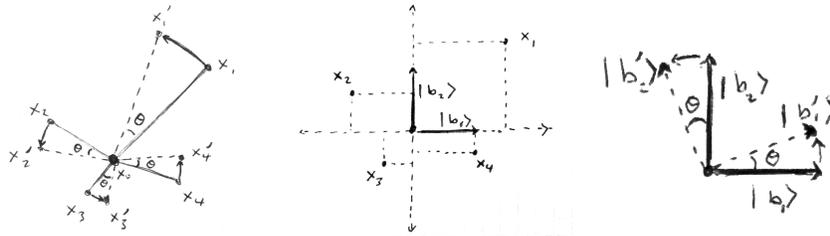
Figure 2: Left: Physical rotation around a point on a wall. Middle: A choice of coordinate system defined by basis $\{|\boldsymbol{b}_1\rangle, |\boldsymbol{b}_2\rangle\}$. Right: Rotation of the basis elements.

Consider the traditional orthogonal,[3] vertically and horizontally aligned, basis shown in the middle plot of Figure 2. Our transform $\boldsymbol{T}$ rotates all points along the wall by an angle θ counterclockwise around the central axis $\boldsymbol{x}_0$; our job now is to find the matrix of this rotation as described in this orthogonal coordinate system. Specifically, $\boldsymbol{T} : \mathcal{W} \to \mathcal{W}$, where $\mathcal{W}$ denotes the 2D vector space we're using to represent the wall, and we've chosen the basis $\mathcal{B} = \{|\boldsymbol{b}_1\rangle, |\boldsymbol{b}_2\rangle\}$ to represent points in both the domain and range spaces of the map (which are both the same space).

So now we have two questions:

1. Is this transform linear?

2. If so, what is the matrix of this transform with respect to the chosen basis $\mathcal{B}$?

For the first question, is this transform linear, we can just go through an intuitive argument to convince ourselves. From the above definitions in Section 2, we know that a linear transform has to satisfy $\boldsymbol{T}\left(\alpha|\boldsymbol{a}\rangle + \beta|\boldsymbol{b}\rangle\right) = \alpha\boldsymbol{T}|\boldsymbol{a}\rangle + \beta\boldsymbol{T}|\boldsymbol{b}\rangle$ (i.e. transforming the linearly combined vector is the same as linearly combining the transformed vectors). In the case of rotations, that would suggest that it doesn't matter whether we add or scale vectors before rotation or after rotation, it all turns out the same. And we know that to be true. Geometry on the wall doesn't change if we tilt our head to the side, it's just rotated, so we can, at least informally for now, convince ourselves that rotation by a fixed angle is a linear operation and proceed to the next question: What is the matrix of this transform in the chosen basis?

From Section 2, we saw that all we need to know to find the matrix representation is the coordinates of each transformed basis element in terms of the range basis. Thus, we need to find the coordinates of the rotated vectors $|\boldsymbol{b}_1'\rangle = \boldsymbol{T}|\boldsymbol{b}_1\rangle$ and $|\boldsymbol{b}_2'\rangle = \boldsymbol{T}|\boldsymbol{b}_2\rangle$ in terms of our chosen orthogonal basis $\mathcal{B}$. These rotated basis elements are shown on the rightmost plot of Figure 2, and we can simply read

---

[3]Technically, we haven't defined inner-products and orthogonality yet, but we're all accustomed to typical Euclidean geometry in the 3D space around us, so I'll assume it for now.

off the coordinates using what we know from trigonometry:

$$|\boldsymbol{b}'_1\rangle \to (\cos\theta, \sin\theta)$$
$$|\boldsymbol{b}'_2\rangle \to (-\sin\theta, \cos\theta).$$

Aligning these elements as the columns of a matrix gives us our matrix representation of the transform in this choice of basis:

$$\boldsymbol{R}(\theta) = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}. \tag{8}$$

Now we can verify that we've done it right by looking at where the basis elements go under this transform in our chosen coordinates. The first basis element $|\boldsymbol{b}_1\rangle$, as represented in $\mathcal{B}$, is just $(1,0)$ (the Kronecker delta). Multiplying by the matrix just extracts the first column, so it just repeats our construction in reverse. The same goes for the second basis element. $|\boldsymbol{b}_2\rangle$, as represented in $\mathcal{B}$, is just $(0,1)$, and multiplying by the matrix $\mathbb{R}(\theta)$ just extracts the second column. Both of these are correctly the coordinates of the rotated points.

References

Dummit, David S. and Foote, Richard M. *Abstract Algebra*. Wiley, 3rd edition, 2003.

Hassani, Sadri. *Mathematical Physics: A Modern Introduction to its Foundations*. Springer, 2nd edition, 2013.

Wikipedia. Bra-ket notation, 2002-2014. URL `http://en.wikipedia.org/wiki/Bra-ket_notation`.