

Vector Visualizations

Dan Calderone

May 28, 2023

Abstract

This tutorial paper gives basic visualization techniques for vectors and inner products. Two different techniques for vector visualization are detailed: *orthogonal (or spatial) axis* representation and *parallel axis* representation. Examples and explanation are given in two, three, and higher dimensions and limitations of each visualization technique are discussed. The geometry of basic sets of vectors are illustrated and discussed and then two strategies for visualizing inner products are given. The first makes use of only the orthogonal representation of vectors and extends the traditional notion of an inner product as a projection to cases where neither vector has unit norm. The second makes use of both the parallel and orthogonal representations and can be easily extended to higher dimensions. ¹

1 Introduction

Vector visualization is at the heart of developing spatial intuition for linear algebra and many other modern topics in modern math and engineering. This paper focuses on techniques for visualizing vectors and inner products.

Two vector visualization strategies are discussed: *orthogonal (or spatial)* representation and *parallel* representation. In the orthogonal representation strategies, vector coordinate values are drawn relative to a coordinate system taken to be orthogonal (in some n -dimensional space). The vector is then visualized relative to a 2D projection of that higher dimensional coordinate system. In the parallel representation, vector coordinate values are drawn along parallel axes.

Basic sets of vectors (unit balls, unit cubes, and simplices) are discussed along with the problem of depth in the orthogonal representation. Two techniques for visualizing inner products are then detailed. The first makes use solely of orthogonal vector representation and extends the intuition behind orthogonal projection to general inner products (where neither vector is a unit vector). The second technique combines the parallel and spatial representations to give a very general technique for visualizing inner products. Variations of this second technique are detailed thoroughly along with their relationship to matrix multiplication and application in higher dimensional contexts.

2 Vector Visualizations

Finite dimensional vectors are represented as a string of digits, each of which gives a displacement relative to an axis. Visual representations of vectors show these displacements in various ways. We will focus on two methods which we will refer to as the *orthogonal (or spatial) axis* representation and the *parallel axis* representation.

¹©Dan Calderone, September 2022

2.1 Spatial (Orthogonal) Axes

The natural (spatial) way to represent the various displacements is along axes that are orthogonal to each other visualized in 2D and 3D here.




Figure 1: Vectors in the spatial representation in \mathbb{R}^2 and \mathbb{R}^3 . Axes are represented by distinct (2D) directions (in the image). We conceptualize these vectors as orthogonal in the higher dimensional space even though the image is a 2D projection.

Cubes (or rectangles) can be used to visualize each individual coordinate of the vector. The length of each cube edge shows the size of each coordinate. Note also that we can visualize building up a vector one coordinate at a time as walking along edges of the cube from the origin to the tip. Note there are multiple different paths that all result in reaching the tip corresponding to the fact that coordinates can be added in any order.




Figure 2: Paths along edges of the hypercube (in the orthogonal representation define different ways the individual coordinates can be added to reach the full vector).

Our brains are highly adapted for visualizing vectors orthogonally in two and three dimensions. If we wish to use the orthogonal representation for vectors in dimensions higher than three, the best we can do is to draw a 2D projection of the higher dimensional vectors. One projection method can be achieved by simply drawing a direction (in 2D) for each axis and then simply showing displacements along these axes. This process is illustrated here for a vector in \mathbb{R}^4 .




Figure 3: Orthogonal representation in \mathbb{R}^4 .

Since there is only room for two orthogonal axes in 2D, we must imagine that the axes we draw are actually orthogonal

Remark 1. *The following exercise is useful for visualizing hypercubes in higher dimension. A 1D cube is a basic line segment along the interval $[0, 1]$. We can obtain a 2D cube, a square, by sweeping that interval along a 2nd axis. Sweeping the square along a 3rd axis produces a 3D cube.*




Figure 4: (Left) A segment is a 1D hypercube. (Center) Sweeping the segment along the second axis gives a square (a 2D hypercube). (Left) Sweeping the square along the third axis gives a cube (a 3D hypercube).

Higher dimensional cubes can be produced by continuing this process. A 3D cube swept along a 4th axis gives a 4D-hypercube; a 4D-hypercube swept along a 5th axis produces a 5D-hypercube; etc.




Figure 5: (Left) Sweeping the cube along the fourth axis gives a 4D-hypercube. (Right) Sweeping the 4D-hypercube along the fifth axis gives a 5D-hypercube.

Remark 2. *When projecting higher dimensional shapes onto 2D images, a certain amount of information, the “depth” direction(s), in the image get lost. If we’re drawing a 3D vector, depth is one dimensional (out of the page). If we’re drawing a 4D vector, depth is 2-dimensional; for a 5D vector, depth is 3-dimensional, etc. Any intuition derived from projections of higher dimensional sets should be verified with rigorous proof.*

Referring to these visualization techniques as an “orthogonal” representation is somewhat of a misnomer because the axes are only truly orthogonal in 2D. We use this term to reference the fact that we are conceptualizing the axes as orthogonal in some higher dimensional space even though we can only view 2D projections. The term *spatial* representation is perhaps more accurate.

2.2 Parallel Axes

Another less traditional way to represent vectors is to place the axes parallel to each other and show each coordinate displacement in the same direction. We will often use rectangles of the appropriate heights to visualize this type of representation as demonstrated here for vectors in 2D, 3D, and 4D.




Figure 6: Parallel representation of vectors in 2D, 3D, and 4D.

Note that for negative coordinates the rectangles extend down from the zero level. Other options are possible particularly visualizing coordinates as points on parallel axes or on the same axis.

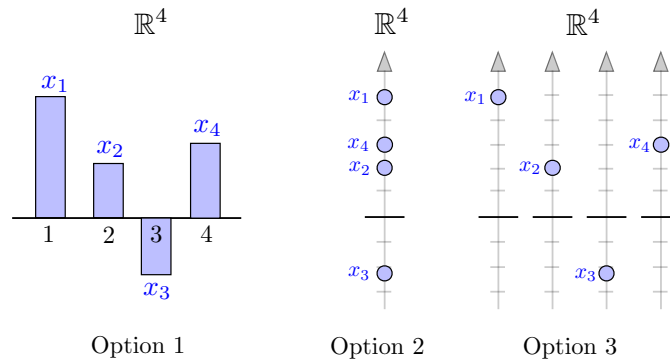


Figure 7: Alternative options for parallel representations of vectors (in 4D).

This parallel axis representation does not suffer from the pathology of depth; however this is simply because it does not seek to leverage our 3D spatial intuition. For example, even in 2D or 3D it is quite difficult to see immediately that two vectors are orthogonal in their parallel axes representation. (Readers are encouraged to try this.) This parallel axis representation will arise in several ways in our treatment of these subjects. First, when we represent columns of a matrix using a spatial representation, the rows naturally appear in a parallel representation along each axis (and vice versa if the rows are represented spatially, the columns appear in a parallel representation.) While not used significantly in this paper, this is an interesting fact that the authors hope to explore in future work. Second, a hybrid parallel-spatial geometric technique can be quite useful in visualizing inner products (discussed below).

3 Basic Geometric Sets

We briefly discuss several geometric sets that we will refer to later. For each set, we label several examples of interesting points. We have already discussed unit cubes briefly. Similarly, let \square_n represent the n -dimensional unit cube.

$$\square_n = \left\{ x \in \mathbb{R}^n \mid \mathbf{0} \leq x \leq \mathbf{1} \right\}$$




Figure 8: Hypercubes in 2D, 3D, and 4D (orthogonal representations).

Let \mathcal{O}_n represent the n -dimensional unit sphere.

$$\mathcal{O}_n = \left\{ x \in \mathbb{R}^n \mid \|x\|_2 = 1 \right\}$$




Figure 9: Hyperspheres in 2D, 3D, and 4D (orthogonal representations).

Note that depending on the axes are drawn a unit circle may not appear perfectly circular. Let Δ_n represent the n -dimensional simplex.

$$\Delta_n = \left\{ x \in \mathbb{R}^n \mid 1^T x = 1, x \geq 0 \right\}$$




Figure 10: Simplices in 2D, 3D, and 4D (orthogonal representations).

Simplexes are often used to represent probability distributions on finite sets, ie. a vector $x \in \mathbb{R}^n$ can represent a discrete probability distribution if and only if it satisfies the above conditions. We note that if we remove the constraints that each element $x_i \geq 0$, we get a particular $n - 1$ -dimensional affine space

$$\ell_n = \left\{ x \in \mathbb{R}^n \mid 1^T x = 1 \right\}$$

For $n = 2$ this is a line through the points I_1, I_2 ; for $n = 3$ this is a plane through the points I_1, I_2, I_3 . Later on, this set will be useful in defining many affine spaces via transformation through a matrix.




Figure 11: Affine sets ℓ_2 and ℓ_3

It will also be useful to represent unit balls of other norms. The 2-norm unit ball is simply the Euclidean unit sphere above. Let \diamond_n represent the n -dimensional unit 1-norm ball.

$$\diamond_n = \left\{ x \in \mathbb{R}^n \mid \|x\|_1 = 1 \right\}$$

Note that the “faces” of the 1-norm ball are simplices with different sign patterns.




Figure 12: 1-norm balls in 2D, 3D, and 4D (orthogonal representation)

Finally, let \square_n^∞ represent the n -dimensional unit ∞ -norm ball.

$$\square_n^\infty = \left\{ x \in \mathbb{R}^n \mid -\mathbf{1} \leq x \leq \mathbf{1} \right\}$$




Figure 13: Inf-norm balls in 2D, 3D, and 4D (orthogonal representation)

Note that the unit ∞ -ball is substantially larger than the unit cube though it has the same shape. The unit cube always has unit volume regardless of the dimension. The volume of the unit ∞ -ball grows rapidly as 2^n .

Again discussion of these sets is left till later.

4 Inner Products

Visualizing inner products, $y^T x$ for $x, y \in \mathbb{R}^n$ is critical to any geometric visualization. We detail two methods here: one based on the spatial axis representations of vectors and one based on a hybrid parallel-spatial axis representation. Note that while we will present these techniques with the two vectors y and x playing different roles, inner products are symmetric and the roles of these vectors can always be reversed.

4.1 Spatial Visualization

The first technique is a somewhat traditional approach to inner product visualization where an inner product is thought of as a projection from one vector onto another. Consider two vectors $x, y \in \mathbb{R}^n$. First define the unit vector in the x direction (labeled u) and an $(n - 1)$ -dimensional hyperplane, \mathcal{T} tangent to the unit sphere at this point (shown here in 2D and 3D). Next, define a vector/point v that points in the y -direction on this plane, given by the intersection of this hyperplane with the line through y . If we drag this point v to the tip of x (and move the tip of y in a parallel motion), then y moves to a vector with length $y^T x$. This process is illustrated (for $y, x \in \mathbb{R}^2$) here.




Figure 14: Spatial visualizations of inner products in 2D (left) and 3D (right) for general x and y

A brief algebraic justification for this is warranted.

Proof. From properties of similar triangles and the definition of v we can obtain that

$$\frac{\lambda}{\|x\|_2} = \frac{\|y\|_2}{\|v\|_2}, \quad \|v\|_2 \cos \theta = 1$$

It follows that $\lambda = \|y\|_2 \|x\|_2 \cos \theta = y^T x$. If x is a unit vector, i.e. $\|x\|_2 = 1$, this visualization becomes the traditional visualization (shown below) of an inner product as a projection. One can think of this technique as extending the projection idea to the case when x and y are both non-unit vectors.




Figure 15: Spatial visualizations of inner products in 2D (left) and 3D (right) when x is a unit vector.

The reader is encouraged to experiment with other values of x and y . Here we show the case where the inner product is negative and the important limiting case where y and x are orthogonal ($y^T x = 0$) and the point v goes to ∞ , i.e. y does not intersect the tangent plane.




Figure 16: Case where $y^T x < 0$ (left) and $y^T x = 0$ (right).

We also note that this technique reduces to perhaps the most natural way to visualize scalar multiplication. For two values $x, y \in \mathbb{R}$ on a number line, one can think of taking the product yx as stretching a unit value for y to be the number x , ie. treating x as the “units” for y . This “rescaling” of the y number line (by x) moves y to yx .




Figure 17: Spatial representation of inner products for scalar y, x .

4.2 Parallel-Spatial

The second inner product visualization technique combines the parallel axis and spatial axis representations of vectors. This technique is subtle and actually quite powerful and is closely related to the idea of column geometry visualization which is at the heart of this paper.

For this technique, we will represent one vector, y , in it’s parallel representation, and the other vector, x , in it’s spatial representation. We illustrate this for $y, x \in \mathbb{R}^3$. First, draw y as a set of heights (or parallel displacements). Next, draw the axes for x with the tips of the standard basis vectors at the ends of each coordinate values of y and place the origin anywhere along the zero-value line. Finally, draw the spatial representation of x relative to these axis. The height of the corresponding point (or, more generally, the displacement along the parallel axis direction) is $y^T x$. This technique is illustrated in the figure below.

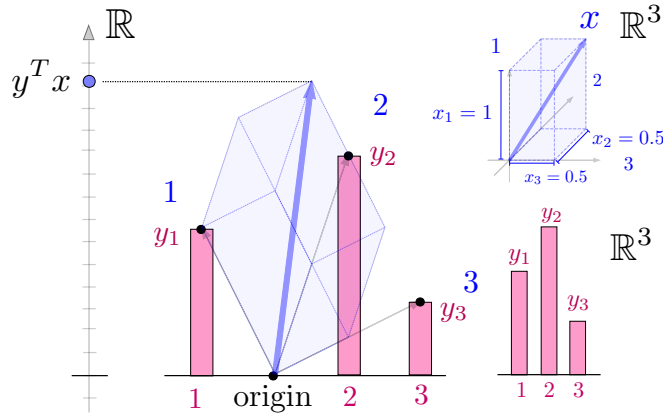


Figure 18: Hybrid parallel-spatial inner-product visualization. x is drawn in it's orthogonal representation relative to axes whose basis vectors run from an (arbitrary) original point on the baseline to the elements of y drawn in parallel representation. The height of x represented in this way is $y^\top x$.

Algebraically, the justification for this technique is immediate.

Proof. For $y, x \in \mathbb{R}^n$, let $Y \in \mathbb{R}^{2 \times n}$ be a matrix whose columns are the 2D vectors that represent the directions of the x -axes drawn on the parallel representation of y , ie. let

$$Y = \begin{bmatrix} -h^\top & - \\ -y^\top & - \end{bmatrix} = \begin{bmatrix} h_1 & h_2 & \cdots & h_n \\ y_1 & y_2 & \cdots & y_n \end{bmatrix}$$

where $h \in \mathbb{R}^n$ is vector of arbitrary horizontal lengths for the axes. For a vector $x \in \mathbb{R}^n$, Yx gives the location of x in the hybrid diagram. Note that

$$Yx = \begin{bmatrix} h^\top x \\ y^\top x \end{bmatrix} = \begin{bmatrix} \sum_i h_i x_i \\ \sum_i y_i x_i \end{bmatrix}$$

and the second coordinate (the height) gives the inner product $y^\top x$. Note the direction of the “second axis” in this example is arbitrary and could point vertically, horizontally or in any other direction. What is more important is that the values y_i are all measured from the same zero line. \square

A more geometric intuition for why this works comes from considering the geometry of parallelograms and similar triangles. We illustrate this here for two dimensions and then extend the argument to three dimensions; one can see how an inductive argument would easily extend it to n -dimensions.




Figure 19: Geometry of the hybrid visualization method starting in 2D (left) and then extending to 3D (right). Note the similar triangles and repeated application of parallelogram geometry.

Key to this idea is that the location of the origin and actually all the horizontal values in this picture are irrelevant to the final result and can be shifted freely. We will often take advantage of this to improve the legibility of various illustrations. Again, the reader is encouraged to try this.

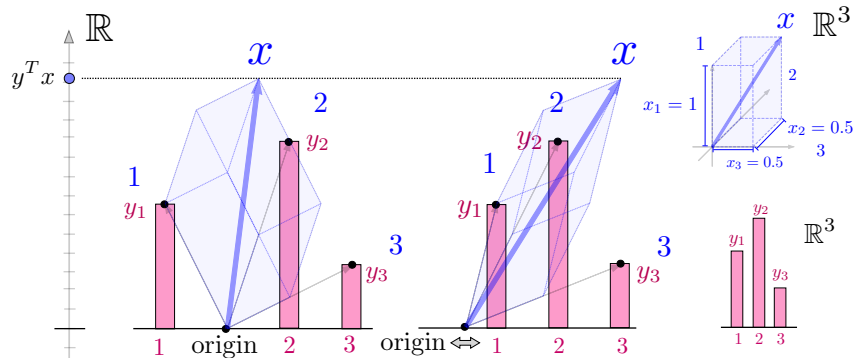


Figure 20: Location of the origin is arbitrary, ie. it does not affect the height and thus the value of $y^T x$.

This technique is particularly useful when the vector x is a convex combination, ie. $\mathbf{1}^T x = 1$, $x \geq 0$. In this case, the image of x in the hybrid visualization will appear in the convex hull of the heights of the y_i 's as illustrated here. This image and then the value $y^T x$ are easy to see.




Figure 21: The method is particularly suited to visualizing $y^\top x$ when x is contained in the simplex, i.e. x takes a convex combination of the values of y .

On the other hand, the visualization can be more difficult to use when the values of x are large or negative. It still works, however.




Figure 22: Hybrid representation when some values of x and/or y are negative.

Other shapes besides convex hulls such as spheres, cubes, etc. are possible as well though (perhaps) more cumbersome.

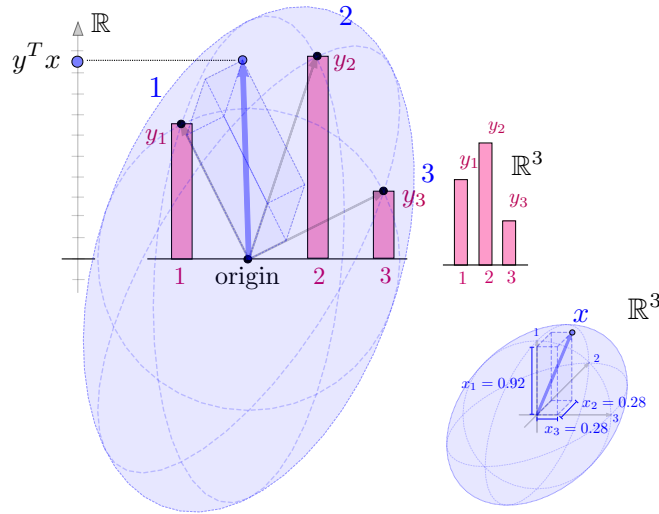


Figure 23: Values of $y^\top x$ for x on the unit sphere.

Visualizing the unit sphere in this way can be especially useful for seeing the set of vectors that are orthogonal to y when y is represented in the parallel representation. The set of vectors orthogonal to x is the preimage of the line at zero height in the domain, the x space shown here.

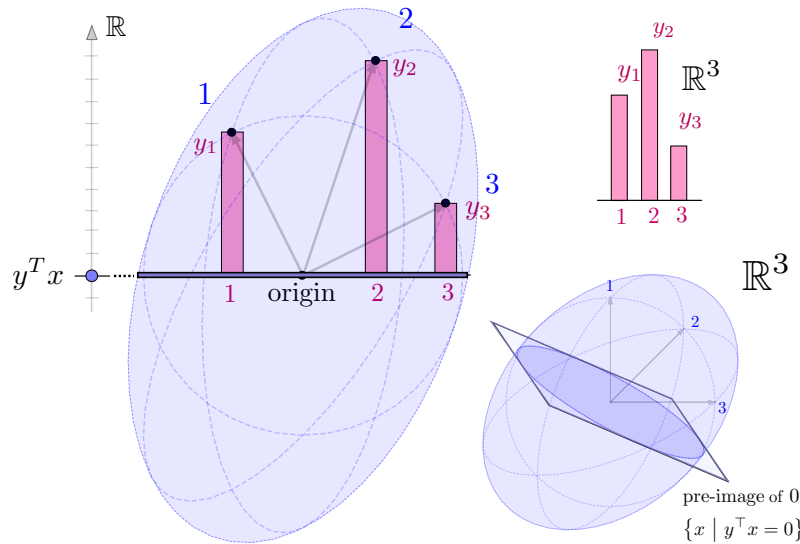


Figure 24: The pre-image of the zero line $y^\top x = 0$ gives the set of vectors x orthogonal to the to y .

It is worth noting that the worst option for the horizontal position of the elements of y is to line them all up on the same axis. Providing variation in the horizontal configuration allows us to differentiate the vector elements and “see” the spatial representation of x better. In fact, many times even if the points naturally lie on the same axis, we will simply “bump” each of them off the axis an arbitrary (but different) amount. When we project the result back onto the original axis to visualize $y^\top x$ we are removing the horizontal information that we added initially.




Figure 25: The horizontal components of the axes are arbitrary and can be adjusted to make the geometry of x more visible.

Also it should be noted that there is no special about “horizontal” and “vertical” in these examples and their roles can be swapped with the same intuition. However, it is important that the “perturbation” direction is orthogonal to the primary direction.

This parallel-spatial inner-product visualization technique works for higher dimensional vectors as well and in this way is more versatile than the classical spatial visualization technique. We demonstrate it here for $x, y \in \mathbb{R}^5$ to for a given y shown in parallel representation and x 's in the simplex and unit cube




Figure 26: Examples of hybrid representation for $x, y \in \mathbb{R}^5$ for x in the simplex (left) and x in the unit cube (right)

and also the unit sphere.

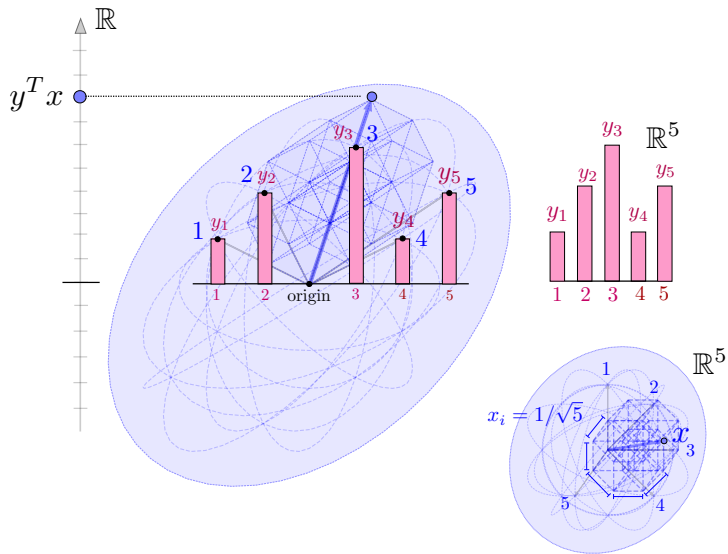


Figure 27: Example of hybrid representation for $x, y \in \mathbb{R}^5$ for x in the unit sphere.

As the dimension increases, one's ability to accurately visualize shapes decreases. However, the spatial intuition can still be useful to get a sense of what the inner product will be.