

Linear Algebra In A Jiffy

Leo Livshits

Tuesday 2nd April, 2024

Dedicated to Heydar Radjavi on the occasion of his 89-th birthday.

Abstract

Every n -tuple in $\mathbb{F}^n$ has a first non-zero entry and a last non-zero entry. What do the positions of such entries in the elements of a subspace $\mathcal{W}$ of $\mathbb{F}^n$ reveal about $\mathcal{W}$? It turns out, a great deal! This insight offers a beeline to the fundamental results on general bases and dimension (including dimensions of complements, rank of a transpose, Rank-Nullity Theorem, Full Rank Factorization, and existence/uniqueness of RREF), bypassing some of the traditional stumbling blocks and time sinks.

Introduction

This article is about doing old things in a new way. It is aimed at mathematicians who teach linear algebra, and it is laid out in a fairly formal way, for efficiency sake.¹ The reader who decides to embark on this trip is advised to exercise a touch of amnesia; otherwise the tenets of the standard approach will not let go of the mind.

Modern developments suggest that initial focus in introductory linear algebra should be placed on $\mathbb{R}^n$ and $\mathbb{C}^n$, with concepts of orthogonality and least squares playing an earlier and more central role, and Singular Value Decomposition done justice within the first course. To be able to implement this without skipping arguments that justify claims, or replacing proper development with commandments, something has to go overboard.

Our approach establishes basic fundamental results quickly, bypassing traditional stumbling blocks and time sinks. The gained efficiency is due to an early introduction of a combinatorial concept of right/left dimension of a subspace, requiring no prior exposure to linear independence, linear systems or row-reduction.

In a course, the material that should precede our point of departure, would include an introduction to the operations on $\mathbb{R}^n$ (or $\mathbb{C}^n$, or $\mathbb{F}^n$), the concept of a subspace, and the fact that spans and ortho-complements generate subspaces. No more than that.

The article is organized as follows. In section 2 we introduce right/left-standard indices. This is a new concept that serves as a trailhead for the shortcut. Section 3 deals with the right/left-basic elements of subspaces and canonical coordinate systems. Section 4 treats dimensionality. In section 5 we connect the canonical coordinate systems of a subspace with those of its ortho-complement.

¹It is a useful feature of the arguments that they are short and are easily converted into a concrete and visual form that students find pleasantly comprehensible and convincing, despite the absence of full generality and rigorous formalism. This is how the material is presented in our forthcoming textbook “Not Your Grandpa’s First Introduction To Linear Algebra.”

The material in sections 6 and 7 assumes some rudimentary familiarity with matrix algebra, including transposition. Again, only the basics are required. In these sections we apply the new methods to establish the equality of row and column ranks, Rank-Nullity theorem, full rank factorization, and existence/uniqueness of the RREF, all as immediate by-products. All of this happens quickly and at a low cost.

Section 8 describes a recursive form of Gauss-Jordan elimination reminiscent of Gram-Schmidt process, that can be used to construct the RREF of a matrix.

Section 9 of the article deals with a combinatorial question that is not usually a part of the course. If one were to mark all of the positions where the elements of a subspace have their first non-zero entry and their last non-zero entry, what configurations can be thus produced?

Throughout this article, $\mathbb{F}$ represents an arbitrary field. We interpret the vectors of $\mathbb{F}^n$ as lists, and as such these can be presented vertically, horizontally, or even diagonally, as long as the ordering is clear. The entries of a vector in $\mathbb{F}^n$ appear in its n **positions**, indexed 1 through n . The i -th entry of $X \in \mathbb{F}^n$ is the element of $\mathbb{F}$ that appears in the i -th position of X , and it is denoted by $X(i)$. The **k -th standard basis vector E_k** in $\mathbb{F}^n$ is the list that has 1 as its k -th entry and zeros in all other positions.

The last non-zero entry of a vector in $\mathbb{F}^n$ is said to be the **terminating entry** of that vector. Obviously the zero vector has none such. If the terminating entry of $X \in \mathbb{F}^n$ occurs in the j -th position, we say that X **terminates** in the j -th position. Similarly the first non-zero entry of an element of $\mathbb{F}^n$ is its **originating entry**, and X **originates** in the j -th position if its originating entry occurs in that position.

2 Right/Left-Standard Indices

Let $\mathcal{W}$ be a subspace of $\mathbb{F}^n$. If $\mathcal{W}$ contains an element that terminates in the j -th position, we say that j is a **right-standard index** of $\mathcal{W}$, and the j -th position is a **right-standard position** for $\mathcal{W}$. For example, if $\mathcal{W}$ contains an element with no zero entries, then the n -th position is a right-standard position for $\mathcal{W}$. Since subspaces are closed under scaling, a position is right-standard for $\mathcal{W}$ if and only if $\mathcal{W}$ contains an element which terminates with a 1 in that position.

To avoid tongue-twisting, we will trade in the term “right-standard” for “**red**”. Why **red**? Well, both terms start with an **r**, “**red**” is a short word, and **red** entries are bound to leave **red** residue, marking the **red** positions.

Similarly, if $\mathcal{W}$ contains an element that originates in the j -th position, then j is a **left-standard index** of $\mathcal{W}$, and the j -th position is a **left-standard position** for $\mathcal{W}$. We write “**lime**” for “left-standard”.

Since results involving **lime** indices, entries and positions are mostly analogous to those involving the **red** versions thereof, we will just focus on the latter, unless we specify otherwise.

The number of **red** positions for $\mathcal{W}$ is said to be the **red-dimension** of $\mathcal{W}$. Obviously the **red-dimension** of the trivial subspace $\{O_n\}$ is zero. We choose this (temporary) strange terminology to highlight the fact that we are not cheating by somehow sneaking in the usual notion of dimension before the underlying concepts and results are established.

Since any non-zero element of $\mathcal{W}$ terminates in a **red** position,

the only element of $\mathcal{W}$ that has zeros in all **red** positions of $\mathcal{W}$ is the zero vector. (*)

As $\mathcal{W}$ is closed under subtraction, it follows that

two elements of $\mathcal{W}$ are equal exactly when they coincide in all **red** positions of $\mathcal{W}$. (**)

In particular, if i is the *smallest* **red** index of $\mathcal{W}$, then there is exactly one element in $\mathcal{W}$ which terminates with a 1 in the i -th position.

The existence of the **red** indices for a non-trivial subspace of $\mathbb{F}^n$ is undeniable, *no matter how the subspace is defined*. Finding what these indices are in a particular situation is a different story, as is often the case in linear algebra. This does not diminish the value of the approach, which is designed to beeline through the introductory fundamentals of linear algebra in an honest way.

3 Red-Basic Elements

3.1 Theorem Red-Basic Elements

When $\mathcal{W}$ is a non-trivial subspace of $\mathbb{F}^n$, for each **red** index i of $\mathcal{W}$ there exists a unique element $W_i \in \mathcal{W}$ that terminates with a 1 in the i -th position, and has 0's in all preceding **red** positions of $\mathcal{W}$.

Proof. Let us start with the uniqueness, as this is an easier task. If X and Y were distinct elements of $\mathcal{W}$ with the described property for a **red** index i , their difference would be a non-zero element of $\mathcal{W}$ that has zeros in all **red** positions of $\mathcal{W}$, which is not possible, by (*).

On to the existence. As we have already observed, the claim holds true for the smallest **red** index i_1 of $\mathcal{W}$. We proceed inductively. Suppose that $W_{i_1}, W_{i_2}, \dots, W_{i_k}$ have been constructed for the smallest k **red** indices $i_1, i_2, \dots, i_k$ of $\mathcal{W}$, and i_{k+1} is the next **red** index of $\mathcal{W}$.

$\mathcal{W}$ contains an element X which terminates with a 1 in the i_{k+1} -st position. It follows that

$$X - X(i_1)W_{i_1} - X(i_2)W_{i_2} - \dots - X(i_k)W_{i_k}$$

is an element of $\mathcal{W}$ that has zero entries in positions indexed by $i_1, i_2, \dots, i_k$, and terminates with a 1 in the i_{k+1} -st position. This is the $W_{i_{k+1}}$ we seek. ■

For each **red** index i of $\mathcal{W}$, we will refer to the element W_i described in theorem 3.1 as **the red-basic element** of $\mathcal{W}$ corresponding to i . Restricted to the **red** positions of $\mathcal{W}$ the **red**-basic elements of $\mathcal{W}$ look like the standard basis m -tuples $E_1, E_2, \dots, E_m$, where m is the **red**-imension of $\mathcal{W}$.

3.2 Theorem Red Entries Are Free

Suppose that a subspace $\mathcal{W}$ of $\mathbb{F}^n$ has red-dimension k , and $a_1, a_2, \dots, a_k \in \mathbb{F}$. Then there exists exactly one element $X \in \mathcal{W}$ with $a_1, a_2, \dots, a_k$ as its entries in the red positions of $\mathcal{W}$ (in that order).

Proof. We have already noted in (**) that there cannot be two distinct elements of $\mathcal{W}$ with the described property. All we need to show is that there is always at least one such.

This is not difficult, since on the red positions of $\mathcal{W}$ the red-basic elements of $\mathcal{W}$ look like the standard basis vectors $E_1, E_2, \dots, E_k$ of $\mathbb{F}^k$.

If $i_1 < i_2 < \dots < i_k$ are the red indices of $\mathcal{W}$, then the element

$$X = a_1 W_{i_1} + a_2 W_{i_2} + \dots + a_k W_{i_k} \in \mathcal{W}$$

satisfies the required condition of having $a_1, a_2, \dots, a_k$ as its entries in the red positions of $\mathcal{W}$. ■

3.3 Definition Coordinate Systems

We say that a list $X_1, X_2, \dots, X_k$ in a subspace $\mathcal{W}$ of $\mathbb{F}^n$ generates a **coordinate system** for $\mathcal{W}$ if every element of $\mathcal{W}$ can be expressed uniquely as a linear combination of $X_1, X_2, \dots, X_k$.

3.4 Corollary Red-basic elements generate a coordinate system

If $\mathcal{W}$ is a non-trivial subspace of $\mathbb{F}^n$ then the red-basic elements of $\mathcal{W}$ generate a coordinate system of $\mathcal{W}$.

Proof. We use the fact that on the red positions of $\mathcal{W}$ the red-basic elements look like the standard basis vectors.

If $i_1 < i_2 < \dots < i_k$ are the red indices of $\mathcal{W}$, and $X \in \mathcal{W}$ then

$$X(i_1) W_{i_1} + X(i_2) W_{i_2} + \dots + X(i_k) W_{i_k}$$

is the unique element of $\mathcal{W}$ with $X(i_1), X(i_2), \dots, X(i_k)$ as its entries in the red positions of $\mathcal{W}$. Of course X itself has these same entries in those same positions. By (**),

$$X(i_1) W_{i_1} + X(i_2) W_{i_2} + \dots + X(i_k) W_{i_k}$$

is the one and only linear combination of the red-basic elements of $\mathcal{W}$ that equals X . ■

To express an element X in $\mathcal{W}$ as a linear combination of the red-basic elements of $\mathcal{W}$, one uses the entries of X in the red positions of $\mathcal{W}$ as coefficients.

In particular, if $W_{i_1}, W_{i_2}, \dots, W_{i_k}$ are the red-basic elements of $\mathcal{W}$, testing whether a given n -tuple X is in $\mathcal{W}$ is as easy as pie: simply test the equality

$$X = X(i_1) W_{i_1} + \dots + X(i_k) W_{i_k}. \quad (1)$$

4 Red Indices Give Dimension

The following result is standard and has an easy enough standard proof, which we omit for the sake of brevity.

4.1 Theorem *Coordinate system criterion*

A list $X_1, X_2, \dots, X_k$ in a subspace $\mathcal{W}$ of $\mathbb{F}^n$ generates a coordinate system for $\mathcal{W}$ (see 3.3) exactly when it spans $\mathcal{W}$, does not start with a zero vector, and has no elements that are linear combinations of the preceding elements.

In view of theorem 4.1, we can use the traditional term “**basis**” and our term “coordinate system” interchangeably. The **red basis** of $\mathcal{W}$ is the coordinate system formed by its **red**-basic elements listed in the order of the corresponding **red** indices (see corollary 3.4).

4.2 Theorem *Monotonicity of red-dimension*

If $\mathcal{W} \preceq \mathcal{V} \preceq \mathbb{F}^n$, then $\text{red-dimension}(\mathcal{W}) \leq \text{red-dimension}(\mathcal{V})$, and the equality of **red**-dimensions holds only when $\mathcal{W} = \mathcal{V}$.

Proof. The **red** indices of a subspace are determined by the terminal positions of its elements. Through enlarging the subspace one can only add to the collection of such indices.

If no new **red** indices have been added in the process, then the **red**-basic elements of $\mathcal{W}$ must be exactly the **red**-basic elements of $\mathcal{V}$, which means that $\mathcal{W}$ and $\mathcal{V}$ have the same **red** basis and are therefore equal. ■

4.3 Theorem *Step-up bound for red-dimension*

For any $X_1, X_2, \dots, X_k, Y \in \mathbb{F}^n$,

$$\begin{aligned} \text{red-dimension}\left(\text{Span}(X_1, X_2, \dots, X_k, Y)\right) \\ \leq \text{red-dimension}\left(\text{Span}(X_1, X_2, \dots, X_k)\right) + 1. \end{aligned}$$

Proof. Let us denote the span of $X_1, X_2, \dots, X_k$ by $\mathcal{W}$, and the span of $X_1, X_2, \dots, X_k, Y$ by $\mathcal{V}$. Elements of $\mathcal{V}$ have the form $Z + \alpha Y$, where $Z \in \mathcal{W}$, $\alpha \in \mathbb{F}$.

We need to demonstrate that we could not have created two new **red** indices for the span, by attaching Y to the original list $X_1, X_2, \dots, X_k$. Had we done so, there would be elements $Z_1 + \alpha Y$ and $Z_2 + \beta Y$ that terminate in distinct non-**red** positions for $\mathcal{W}$ (say, i -th and j -th positions, with $i < j$).

Neither α nor β can be zero. Consequently $\beta Z_1 + \alpha \beta Y$ and $\alpha Z_2 + \alpha \beta Y$ still terminate in the i -th and j -th positions respectively. Therefore their difference, which is $\alpha Z_2 - \beta Z_1$, terminates in the j -th position (since $i < j$). Yet this difference is an element of $\mathcal{W}$, and hence cannot terminate in a non-**red** position of $\mathcal{W}$. This contradiction completes the proof. ■

4.4 Theorem Red-dimension is the dimension

Every basis of $\mathcal{W}$ has exactly the same number of elements as the red basis of $\mathcal{W}$.*

*It is worth reiterating that naturally we have made neither use of nor assumption about the equality of cardinalities of various bases of a given subspace. The equality is an easy byproduct of our tack.

Proof. By theorem 4.1, a list $X_1, X_2, \dots, X_m$ in a non-trivial subspace $\mathcal{W}$ of $\mathbb{F}^n$ is a basis of $\mathcal{W}$ exactly when

$$\{0\} \subsetneq \text{Span}(X_1) \subsetneq \text{Span}(X_1, X_2) \subsetneq \dots \subsetneq \text{Span}(X_1, X_2, \dots, X_m) = \mathcal{W}.$$

By theorem 4.2, this, apart from the requirement that $X_i \in \mathcal{W}$, can be restated in terms of red-dimension :

$$\begin{aligned} 1 = \text{red-dimension}(\text{Span}(X_1)) &< \text{red-dimension}(\text{Span}(X_1, X_2)) < \dots \\ &< \text{red-dimension}(\text{Span}(X_1, X_2, \dots, X_m)) = \text{red-dimension}(\mathcal{W}). \end{aligned}$$

In such a case each step-up in red-dimension must be exactly 1, by theorem 4.3, and the proof is complete. ■

4.5 Corollary

Every subspace $\mathcal{W}$ of $\mathbb{F}^n$ has as many lime indices and red ones.*

*Of course the reader will have noticed that the linear isomorphism on $\mathbb{F}^n$ that reverses the order of the entries in each n-tuple, maps a subspace $\mathcal{W}$ to a subspace $\mathcal{W}^{+p}$ in such a way that the red indices of $\mathcal{W}$ become the lime indices of $\mathcal{W}^{+p}$. This demonstrates that red-dimension and lime-dimension of a subspace are equal, but obviously by "borrowing from the future".

A compelling reason for introducing red/lime indices is to provide an efficient alternate path for the first linear algebra course. The red/lime indices are designed to be encountered well before linear functions make their official appearance, or the traditional concept of dimension is introduced.

5 Red/Lime Bases And •-Complements

We will write $X \bullet Y$ to indicate the standard symmetric bilinear form, commonly known as the dot product, for elements of $\mathbb{F}^n$. For each collection $\mathcal{S}$ in $\mathbb{F}^n$, we define its **•-complement** to be

$$\mathcal{S}^\bullet = \{ X \in \mathbb{F}^n \mid \forall Y \in \mathcal{S} : X \bullet Y = 0 \}.$$

When $\mathbb{F} = \mathbb{R}$, •-complement captures the geometry of perpendicularity. ²

²For the standard inner product on $\mathbb{C}^n$ see remark 5.4.

5.1 Theorem Duality of red/lime indices of •-complements

For any subspace $\mathcal{W}$ of $\mathbb{F}^n$, the lime indices of $\mathcal{W}^\bullet$ are exactly the non-red indices of $\mathcal{W}$.

Proof. If i is a red index of $\mathcal{W}$, it cannot be a lime index of $\mathcal{W}^\bullet$, since otherwise $\mathcal{W}$ would contain an element that terminates with a 1 in the i -th position, while $\mathcal{W}^\bullet$ would contain an element that originates with a 1 in the i -th position. The dot product of these two elements would be 1, contrary the definition of $\mathcal{W}^\bullet$.

If i_o is a non-red index of $\mathcal{W}$, then there are two scenarios to consider. In the first scenario there are no red indices of $\mathcal{W}$ that are greater than i_o . In that case every element of $\mathcal{W}$ terminates before the i_o -th position. Therefore the standard basis vector E_{i_o} is in $\mathcal{W}^\bullet$, which shows that i_o is a lime index of $\mathcal{W}^\bullet$.

In the remaining scenario i_o is followed by some red indices, say $i_1, i_2, \dots, i_k$. Let us write a_j for the i_o -th entry of W_{i_j} .

We shall construct an element Z of $\mathcal{W}^\bullet$ which originates in the i_o -th position, so that i_o is a lime index of $\mathcal{W}^\bullet$.³ The design of Z is such that for each red-basic element W_i of $\mathcal{W}$, there are at most two positions where both W_i and Z have non-zero entries.

We let

$$Z(m) = \begin{cases} 1 & \text{if } m = i_o \\ -a_j & \text{if } m = i_j; 1 \leq j \leq k \\ 0 & \text{otherwise} \end{cases}.$$

It is easy to see that $Z \bullet W = 0$ for every red-basic element W_i of $\mathcal{W}$, and since these form a basis of $\mathcal{W}$, $Z \in \mathcal{W}^\bullet$. ■

Here is an illustration of the construction in the proof of theorem 5.1, when $i_o = 2$ and $i_1, i_2, \dots, i_k$ are 5, 8, 9, 13, $\dots$, p :

$$\begin{array}{lcl} W_5 & = & (*, a_5, *, *, 1, 0, 0, 0, 0, 0, 0, 0, 0, \dots, 0, \dots, 0) \\ W_8 & = & (*, a_8, *, *, 0, *, *, 1, 0, 0, 0, 0, 0, \dots, 0, \dots, 0) \\ W_9 & = & (*, a_9, *, *, 0, *, *, 0, 1, 0, 0, 0, 0, \dots, 0, \dots, 0) \\ & \vdots & \\ W_p & = & (*, a_p, *, *, 0, *, *, 0, 0, *, *, *, 0, \dots, 1, \dots, 0) \\ Z & = & (0, \underset{\uparrow}{1}, 0, 0, \underset{\uparrow}{-a_5}, 0, 0, \underset{\uparrow}{-a_8}, \underset{\uparrow}{-a_9}, 0, 0, 0, \underset{\uparrow}{-a_{13}}, \dots, \underset{\uparrow}{-a_p}, \dots, 0) \end{array}$$

5.2 Observation

It is worth noting that the element Z constructed in the proof of theorem 5.1 is exactly the lime-basic element of $\mathcal{W}^\bullet$ corresponding to the lime index i_o of $\mathcal{W}^\bullet$. This shows that

the lime basis of $\mathcal{W}^\bullet$ can be simply read off from the red basis of $\mathcal{W}$, and vice versa.

³This proof is followed by an illustration of the construction.

5.3 Corollary Dimensions of •-complements

For any subspace $\mathcal{W}$ of $\mathbb{F}^n$, $\dim(\mathcal{W}) + \dim(\mathcal{W}^\bullet) = n$, and consequently $(\mathcal{W}^\bullet)^\bullet = \mathcal{W}$.

Proof. We have verified that the lime indices of $\mathcal{W}^\bullet$ are exactly the non-red indices of $\mathcal{W}$. Of course this shows that $\dim(\mathcal{W}^\bullet) = n - \dim(\mathcal{W})$.

To establish the remaining claim, note that $\mathcal{W}$ is naturally included in $(\mathcal{W}^\bullet)^\bullet$, and the two subspaces have the same dimensions. ■

5.4 Observation

An analogous argument (cf. proof of theorem 5.1) demonstrates similar results in the case $\mathbb{F} = \mathbb{C}$, when $\bullet$ is replaced by the standard inner product. This can be viewed as a particular case of a more general claim. If $\varphi : \mathbb{F} \rightarrow \mathbb{F}$ is an involutive isomorphism, and we define $\mathcal{P} : \mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}$ by

$$\mathcal{P}(X, Y) = X \bullet \varphi^{(n)}(Y),$$

where

$$\varphi^{(n)} \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} \varphi(y_1) \\ \varphi(y_2) \\ \vdots \\ \varphi(y_n) \end{pmatrix},$$

then analogous conclusions can be drawn regarding $\mathcal{W}^\mathcal{P}$ and $(\mathcal{W}^\mathcal{P})^\mathcal{P}$.

6 Red/Lime Bases And Ranges/Nullspaces

Let us begin with a brief review of some concepts and terminology pertaining to matrices, which would have to be introduced prior to the material we are about to discuss.

When one applies an $n \times m$ matrix $\mathbf{A}$ to an m -tuple X , the resulting n -tuple $\mathbf{A}(X)$ can be described in terms of the rows $R_1, R_2, \dots, R_n$ of $\mathbf{A}$, as well as in terms of the columns $C_1, C_2, \dots, C_m$ of $\mathbf{A}$:

1. the i -th entry of $\mathbf{A}(X)$ is the dot products $R_i \bullet X$;
2. $\mathbf{A}(X) = X(1)C_1 + X(2)C_2 + \dots + X(m)C_m$.

For convenience let us refer to these as a **row-centric** and a **column-centric** descriptions of $\mathbf{A}(X)$, respectively.

The column-centric perspective shows that, interpreted as a linear function acting from $\mathbb{F}^m$ to $\mathbb{F}^n$, $\mathbf{A}$ is invertible (i.e. bijective) if and only if every element of $\mathbb{F}^n$ can be expressed uniquely as a linear combination of the columns of $\mathbf{A}$; in other words, if and only if the columns of $\mathbf{A}$ form basis of $\mathbb{F}^n$ (see definition 3.3 and theorem 4.1). As we have already established in theorem 4.4, every basis of $\mathbb{F}^n$ has n elements. Hence invertible matrices are square.

The row-centric perspective shows that $\mathbf{A}(X)$ is a zero vector exactly when X is in the $\bullet$ -complement of rows of $\mathbf{A}$ within $\mathbb{F}^m$; in other words, when X is in the $\bullet$ -complement of the span of these rows, known as the **row space** of $\mathbf{A}$. This states that

$\text{Nullspace}(\mathbf{A})$ and $\text{Row Space}(\mathbf{A})$ are $\bullet$ -complements within $\mathbb{F}^m$.

While we are on the topic, let us remind the reader that the span of the columns of a matrix $\mathbf{A}$ is said to be its **column space**, and is exactly the range of $\mathbf{A}$, when $\mathbf{A}$ is interpreted as a linear function from $\mathbb{F}^m$ to $\mathbb{F}^n$. **Column rank** of $\mathbf{A}$ is defined to be the dimension of the column space of $\mathbf{A}$, and **row rank** is the dimension of the row space. Since columns of $\mathbf{A}^T$ (the transpose of $\mathbf{A}$) are rows of $\mathbf{A}$, column space of $\mathbf{A}^T$ equals row space of $\mathbf{A}$, and row rank of $\mathbf{A}$ equals column rank of $\mathbf{A}^T$.

By corollary 5.3,

$$m = \text{Nullity}(\mathbf{A}) + \text{Row Rank}(\mathbf{A}) = \text{Nullity}(\mathbf{A}) + \text{Column Rank}(\mathbf{A}^T). \quad (2)$$

Using the column-centric perspective and theorem 3.1, it is not hard to see that each of the following claims regarding an index i and a matrix $\mathbf{A}$ with columns $C_1, C_2, \dots, C_m$ is equivalent to the rest.

1. C_i is a linear combination of the preceding columns of $\mathbf{A}$.
2. There are scalars $a_1, a_2, \dots, a_{i-1} \in \mathbb{F}$ such that

$$0 = a_1 C_1 + a_2 C_2 + \dots + a_{i-1} C_{i-1} + 1 C_i$$

3. There is a vector X in the nullspace of $\mathbf{A}$ that terminates with a 1 in the i -th position.
4. i is a **red** index for the nullspace of $\mathbf{A}$.
5. There is a vector W_i in the nullspace of $\mathbf{A}$ that terminates with a 1 in the i -th position, and has zeros in all of the preceding **red** positions.
6. C_i is a linear combination of the preceding columns of $\mathbf{A}$ that correspond to **non-red** indices of the nullspace of $\mathbf{A}$.

Since $\text{Nullspace}(\mathbf{A})$ and $\text{Row Space}(\mathbf{A})$ are $\bullet$ -complements, the **red** indices of the nullspace of $\mathbf{A}$ are exactly the **non-lime** indices of the row space of $\mathbf{A}$ (by theorem 5.1). It follows, by theorem 4.1, that

if $i_1 < i_2 < \dots < i_k$ are the **lime** (resp. **red**) indices of the row space of $\mathbf{A}$, then the columns $C_{i_1}, C_{i_2}, \dots, C_{i_k}$ form a basis of the column space of $\mathbf{A}$. * (†)

*Of course the roles of rows and columns can be reversed here through transposition.

Therefore

$$\text{Row Rank}(\mathbf{A}) = \text{Column Rank}(\mathbf{A}),$$

and we just use $\text{Rank}(\mathbf{A})$ for both. In particular,

$$\text{Rank}(\mathbf{A}^T) = \text{Rank}(\mathbf{A}). \quad (3)$$

By (2), we have arrived that

$$\text{Rank}(\mathbf{A}) + \text{Nullity}(\mathbf{A}) = m \quad (\text{Rank-Nullity Formula}).$$

7 Red/Lime Bases and RREF/RCEF

Next we discuss the connection between RREF/RCEF and **lime** bases, touching base with Grinberg's "gauche perspective" [1]. Existence and uniqueness of $\text{RREF}(\mathbf{A})$, for any matrix $\mathbf{A}$, is an immediate by-product of our approach.

Our path to RREF/RCEF passes through the Full Rank factorization, which is fundamental in its own right. If the reader has not encountered it previously, we recommend treating it as a generalization of the fact that matrices of rank 1 are exactly those of the form $[\mathbf{X}]_{n \times 1} [\mathbf{Y}]_{m \times 1}^T$.

7.1 Theorem Full Rank factorization using rows of $\mathbf{A}$

Suppose that a matrix $\mathbf{A} \in \mathbb{M}_{n \times m}$ has rank $r > 0$, and $W_{i_1}, W_{i_2}, \dots, W_{i_r}$ is the **lime** basis of the columns space of $\mathbf{A}$.

Let $\mathbf{B}$ be the matrix whose columns are $W_{i_1}, W_{i_2}, \dots, W_{i_r}$, and let $\mathbf{G}$ be the matrix whose rows are the rows $R_{i_1}, R_{i_2}, \dots, R_{i_r}$ of $\mathbf{A}$. * Then

$$\mathbf{A} = \mathbf{B}\mathbf{G} \quad \text{and} \quad \text{Rank}(\mathbf{B}) = r = \text{Rank}(\mathbf{G}).$$

*By (†) these rows form a basis of the row space of $\mathbf{A}$.

Proof. The rank condition is immediate for $\mathbf{B}$, and for $\mathbf{G}$ it follows from (†).

The fact that $\mathbf{A} = \mathbf{B}\mathbf{G}$ is best seen one column at a time, via the column-centric perspective. Each column C_j of $\mathbf{A}$ is a linear combination of $W_{i_1}, W_{i_2}, \dots, W_{i_r}$, i.e. of the columns of $\mathbf{B}$, with coefficients being the entries of C_j corresponding to the indices $i_1 < i_2 < \dots < i_r$ (see (1)). These entries in the columns $C_1, C_2, \dots, C_m$ form rows of $\mathbf{A}$ that correspond to the indices $i_1 < i_2 < \dots < i_r$; i.e. the rows of $\mathbf{G}$. ■

A matrix whose rows are the **lime** basis of its row space, in the increasing order of indices, followed perhaps by some zero rows at the bottom, is said to be in **Reduced Row Echelon Form** (RREF).

If one switches "rows" to "columns" (and "bottom" to "right") in the paragraph above, one arrives at **Reduced Column Echelon Form** (RCEF).

It is not hard to see that the transpose of an RREF matrix is an RCEF matrix, and conversely.

The use **lime** bases renders the question of uniqueness of RREF/RCEF trivial: for a given $m \times n$ matrix $\mathbf{A}$, we write $\text{RREF}(\mathbf{A})$ for the unique $m \times n$ RREF matrix with the same row space as $\mathbf{A}$.

Since row space of a matrix is the $\bullet$ -complement of its nullspace, $\text{RREF}(\mathbf{A})$ is the unique $m \times n$ RREF matrix with the same nullspace as $\mathbf{A}$.⁴ By remark 5.2, one can read off the **red** basis of the nullspace of an RREF matrix from the matrix itself with minimal effort.

Similarly, we write $\text{RCEF}(\mathbf{A})$ for the unique $m \times n$ RCEF matrix with the same column space as $\mathbf{A}$. Since column space of $\mathbf{A}$ equals row space of $\mathbf{A}^T$, the two

⁴In particular, any RREF matrix is completely determined by its nullspace.

spaces share the **lime** basis, and it follows that

$$\left(\text{RCEF}(\mathbf{A})\right)^T = \text{RREF}(\mathbf{A}^T). \quad (4)$$

7.2 Corollary RCEF factorization

Suppose that $\mathbf{A}$ is an $n \times m$ non-zero matrix, and $i_1 < i_2 < \dots < i_r$ are the **lime** indices of the column space of $\mathbf{A}$. Then

$$\mathbf{A} = \text{RCEF}(\mathbf{A}) \circ \mathbf{S}, \quad (5)$$

where $\mathbf{S}$ is the matrix whose rows are the rows $R_{i_1}, R_{i_2}, \dots, R_{i_r}$ of $\mathbf{A}$ (in that order), followed by some zero rows.

By replacing the zero rows at the bottom of $\mathbf{S}$ by some standard basis elements of $\mathbb{F}^m$, one can create an invertible matrix $\mathbf{S}$ which still satisfies (5).

Proof. Take matrices $\mathbf{B}$ and $\mathbf{G}$ described in theorem 7.1. If we beef $\mathbf{B}$ up to an $n \times m$ size, by attaching zero columns from the right, the resulting matrix is exactly $\text{RCEF}(\mathbf{A})$. Append a corresponding number of zero rows to the bottom of $\mathbf{G}$ to produce $\mathbf{S}$.

Let us settle the second claim. The rows of $\mathbf{G}$ form a basis of the row space of $\mathbf{A}$, and so can be extended to a basis of $\mathbb{F}^m$ through addition of some standard basis elements of $\mathbb{F}^m$.⁵ If we replace the zero rows at the bottom of $\mathbf{S}$ with these, we produce the desired new invertible matrix $\mathbf{S}$. ■

If we apply corollary 7.2 to $\mathbf{A}^T$, use equalities (3) and (4), remembering that column space of $\mathbf{A}^T$ and $\text{Nullspace}(\mathbf{A})$ are $\bullet$ -complements, we obtain a corresponding claim about RREF factorization.

7.3 Corollary RREF factorization

Suppose that $\mathbf{A}$ is an $n \times m$ non-zero matrix, and $i_1 < i_2 < \dots < i_r$ are the **lime** indices of the row space of $\mathbf{A}$.^{*} Then

$$\mathbf{A} = \mathbf{T} \circ \text{RREF}(\mathbf{A}), \quad (6)$$

where the columns of $\mathbf{T}$ are the columns $C_{i_1}, C_{i_2}, \dots, C_{i_r}$ of $\mathbf{A}$ (in that order), followed by some zero columns.

By replacing the eastern zero columns of $\mathbf{T}$ by some some standard basis elements of $\mathbb{F}^n$, one can create an invertible matrix $\mathbf{T}$ which still satisfies (6).

^{*}These **lime** indices of the row space of $\mathbf{A}$ are exactly the non-**red** indices of the nullspace of $\mathbf{A}$.

⁵Of course this needs to be demonstrated, but it is an easy argument.

8 Gauss-Jordan à la Gram-Schmidt

The reader may have grown concerned that our approach does not expose students to Gauss-Jordan elimination that has long been sacred in a first course on linear algebra.

We live in the contemporary world with vast computing resources and AI at our fingertips. It is counterproductive to let students think that proficiency in carrying out Gauss-Jordan elimination by hand is a valuable intellectual accomplishment. Proficiency in reducing a quest to a recognizable procedure, being able to use computers to carry out the calculations, and having the skills to interpret the results, is of course a different matter.

The efficiency we gain through our approach is in part due to fact that row-reduction no longer needs to play a central role in the course. Still, Gauss-Jordan elimination has a natural place within the development of *red/lime* bases, *well before matrices are officially introduced*. Furthermore, in this form the elimination is reminiscent of Gram-Schmidt process, or vice versa. This is something that students appreciate.

The value of including this algorithm in a first linear algebra course is in demonstrating its existence. We certainly do not suggest that students should be asked to carry out such a procedure by hand; one should use a computer for any such task.

Like Gram-Schmidt process, the procedure we present is recursive, and it relies on the repeated use of a procedure $\mathfrak{P}$ (described below) which pertains to the following question: when one appends an element $Y \in \mathbb{F}^n$ to the *lime* basis of a subspace $\mathcal{W}$ of $\mathbb{F}^n$, how does one generate the *lime* basis of the resulting span $\mathcal{V}$ of the new list?⁶ We consider the *lime* case because it corresponds most directly to the traditional Gauss-Jordan elimination.

Before stating a general form of the procedure, let us illustrate it with a concrete example. This will convince the reader that the procedure is indeed a form of Gauss-Jordan elimination.

Say, $W_5, W_8, W_9, W_{13}, \dots, W_p$ are the k *lime*-basic elements of $\mathcal{W}$ (that correspond to the *lime* indices $5, 8, 9, 13, \dots, p$ and form the *lime* basis of $\mathcal{W}$). Suppose Y originates with a 1 in the 11-th position, and insert Y into the list between W_9 and W_{13} :

$$\begin{aligned}
 W_5 &= (0, 0, 0, 0, \textcolor{limegreen}{1}, *, *, \textcolor{limegreen}{0}, \textcolor{limegreen}{0}, a, *, *, \textcolor{limegreen}{0}, \dots, \textcolor{limegreen}{0}, \dots, *) \\
 W_8 &= (0, 0, 0, 0, \textcolor{limegreen}{0}, 0, 0, \textcolor{limegreen}{1}, \textcolor{limegreen}{0}, b, *, *, \textcolor{limegreen}{0}, \dots, \textcolor{limegreen}{0}, \dots, *) \\
 W_9 &= (0, 0, 0, 0, \textcolor{limegreen}{0}, 0, 0, \textcolor{limegreen}{0}, \textcolor{limegreen}{1}, c, *, *, \textcolor{limegreen}{0}, \dots, \textcolor{limegreen}{0}, \dots, *) \\
 Y &= (0, 0, 0, 0, \textcolor{limegreen}{0}, 0, 0, \textcolor{limegreen}{0}, \textcolor{limegreen}{0}, 1, *, *, \textcolor{limegreen}{\alpha}, \dots, \textcolor{limegreen}{\zeta}, \dots, *) \\
 W_{13} &= (0, 0, 0, 0, \textcolor{limegreen}{0}, 0, 0, \textcolor{limegreen}{0}, \textcolor{limegreen}{0}, 0, 0, 0, \textcolor{limegreen}{1}, \dots, \textcolor{limegreen}{0}, \dots, *) \\
 &\vdots \\
 W_p &= (0, 0, 0, 0, \textcolor{limegreen}{0}, 0, 0, \textcolor{limegreen}{0}, \textcolor{limegreen}{0}, 0, 0, 0, \textcolor{limegreen}{0}, \dots, \textcolor{limegreen}{1}, \dots, *)
 \end{aligned}$$

⁶Since there is nothing to do if Y is already in $\mathcal{W}$, and testing for this inclusion is easy (see (1)), we will focus on the case of $Y \notin \mathcal{W}$.

Let $\hat{Y} = Y - \alpha W_{13} - \dots - \zeta W_p$. Then

$$\begin{aligned} W_5 &= (0, 0, 0, 0, \textcolor{red}{1}, *, *, 0, 0, a, *, *, 0, \dots, 0, \dots, *) \\ W_8 &= (0, 0, 0, 0, 0, 0, 0, \textcolor{red}{1}, 0, b, *, *, 0, \dots, 0, \dots, *) \\ W_9 &= (0, 0, 0, 0, 0, 0, 0, 0, \textcolor{red}{1}, c, *, *, 0, \dots, 0, \dots, *) \\ \hat{Y} &= (0, 0, 0, 0, 0, 0, 0, 0, 0, 1, *, *, 0, \dots, 0, \dots, *) \\ W_{13} &= (0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, \textcolor{red}{1}, \dots, 0, \dots, *) \\ &\vdots \\ W_p &= (0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, \dots, \textcolor{red}{1}, \dots, *) \end{aligned}$$

In the list $W_5, W_8, W_9, \hat{Y}, W_{13}, \dots, W_p$ replace W_5, W_8, W_9 with

$$W_5 - \alpha \hat{Y}, \quad W_8 - b \hat{Y}, \quad W_9 - c \hat{Y}$$

to produce the **lime** basis of the span of $W_5, W_8, W_9, Y, W_{13}, \dots, W_p$.⁷

Procedure $\mathfrak{P}$: When Y is not in $\mathcal{W}$, by the **lime** version of the **red** case arguments presented on the preceding page, $\mathcal{V}$ gains one new **lime** index in addition to those inherited from $\mathcal{W}$: the originating index i_o of Y . After scaling, we may assume that Y originates with a 1.

By theorem 3.2, some linear combination X of the **lime**-basic elements of $\mathcal{W}$, corresponding to the **lime** indices exceeding i_o , matches Y in every **lime** position of $\mathcal{W}$. Let $\hat{Y} = Y - X$. This $\hat{Y} \in \mathcal{V}$ is not a zero vector and still originates with a 1 in the i_o -th position. The newly minted $\hat{Y}$ is the **lime**-basic element of $\mathcal{V}$ corresponding to the index i_o .

The original **lime**-basic elements of $\mathcal{W}$ corresponding to the **lime** indices that exceed i_o also serve as **lime**-basic elements of $\mathcal{V}$. The **lime**-basic elements of $\mathcal{W}$ corresponding to indices preceding i_o would be **lime**-basic elements of $\mathcal{V}$, were not for their potentially non-zero i_o -th entries. So we subtract from them appropriate multiples of $\hat{Y}$ to zero out these entries, and, voilà, we arrive at the **lime** basis of $\mathcal{V}$.

By using the procedure $\mathfrak{P}$ recursively, one can produce the **lime** basis $Y_1, Y_2, \dots, Y_p$ of the span of the list $0 \neq X_1, X_2, \dots, X_k$ in $\mathbb{F}^n$, with well-aligned intermediate spans, à la Gram-Schmidt process.

Start by scaling X_1 so that it originates with a 1. This is Y_1 , and the case $k = 1$ is done. Then use the procedure $\mathfrak{P}$ recursively on k , by appending X_{j+1} to the already constructed **lime** basis $Y_1, Y_2, \dots, Y_l$ of the span of $X_1, X_2, \dots, X_j$.

Of course this algorithm can be used to construct the **lime** basis of the row space of $\mathbf{A}$, i.e. to construct $\text{RREF}(\mathbf{A})$.

9 Sub-Terminal Entries, Subspace Truncations And Possible Configurations

This section is motivated by a natural question (which turns out to have a pretty answer): “what configurations of **red** and **lime** indices are possible?”

⁷Of course, we could have done the same type of thing with **red** bases.

We will work our way towards the answer by asking a simpler question: “what can be said about the positions of the second-to-last non-zero entry in the elements of a subspace $\mathcal{W}$ of $\mathbb{F}^n$?”

Let us say that $Z \in \mathcal{W}$ **sub-terminates** in the j -th position if its second-to-last non-zero entry is in that position. If an element has no second-to-last non-zero entry, its sub-terminal index is 0.

When X and Y are distinct elements of $\mathcal{W}$ that terminate in the i -th position, after scaling we can assume that their terminating non-zero entries (in the i -th position) are both 1. Hence $X - Y$ is an element of $\mathcal{W}$ that terminates in some **red** k -th position, with $k < i$. In particular, X and Y coincide past the k -th position, but not in the k -th position. Consequently, the sub-terminal index of at least one of X, Y is at least k (and is less than i). If it is strictly between k and i , then it is the sub-terminal index for both X and Y . This shows that when the sub-terminal entries of X and Y occur in non-**red** positions, they occur in the same position.

This shows that

there is at most one non-**red** position in which the elements of $\mathcal{W}$ that terminate in the i -th position can sub-terminate.

The sub-terminal index j_i of a **red**-basic element W_i of $\mathcal{W}$ (3.1) is not **red**, and so it is the only non-**red** sub-terminal index for elements of $\mathcal{W}$ that terminate in the i -th position. Hence we can identify all of the sub-terminal indices of the elements of $\mathcal{W}$ that terminate in the i -th position.

These indices are: j_i , and all **red** indices between j_i and i . *

*To reach this conclusion one makes use of the **red**-basic elements for indices preceding i .

It will be handy to designate the n indices/positions of a subspace $\mathcal{W}$ of $\mathbb{F}^n$ by letters ρ, λ, β and ν as follows:

$$\left\{ \begin{array}{l} \rho : \text{red, but not lime} \\ \lambda : \text{lime, but not red} \\ \beta : \text{both red and lime} \\ \nu : \text{neither red nor lime} \end{array} \right.$$

The last position for a subspace $\mathcal{W}$ of $\mathbb{F}^n$ is a β -position exactly when $\mathcal{W}$ contains the standard basis element E_n .

The last position is a ν -position if and only if all elements of $\mathcal{W}$ have zero as their last entry.

The last position of a subspace can never be a λ -position, and the first position - can never be a ρ -position.

We can transform $\mathbb{F}^n$, with $n > 1$, into $\mathbb{F}^{n-1}$, by simply dropping off the last entry of the n -tuples to produce $(n-1)$ -tuples. In the process subspaces of $\mathbb{F}^n$ get transformed into subspaces of $\mathbb{F}^{n-1}$. For simplicity, let us refer to this as a “truncation of $\mathbb{F}^n$ (and its subspaces) from the right.”

9.1 Theorem Effect of truncation on red/lime indices

The act of a truncation of $\mathcal{W}$ from the right to produce a subspace $\mathcal{U}$ has the following effect on the special positions:

- If the last position of $\mathcal{W}$ is a β -position or a ν -position, then the indices of $\mathcal{U}$ have exactly the same designations as they had for $\mathcal{W}$.
- If the last position of $\mathcal{W}$ is a ρ -position, then one of the indices for $\mathcal{U}$ gains a **red** status (that it did not have for $\mathcal{W}$), and status of all other indices remain unchanged. So, either some ν -index of $\mathcal{W}$ becomes a ρ -index for $\mathcal{U}$, or a λ -index for $\mathcal{W}$ becomes a β -index for $\mathcal{U}$.

Proof. The only possible change in status is that sub-terminal positions of the elements of $\mathcal{W}$ that terminate in the last (n -th) position, become terminal positions for the corresponding elements of $\mathcal{U}$.

As we have observed, such sub-terminal positions are either the sub-terminal position of the **red**-basic element of $\mathcal{W}$ corresponding to the index n (if n happens to be a **red** index) or are the positions corresponding to the **red** indices that are less than n . ■

Now we are ready to answer our original question. Suppose that n positions are marked by (perhaps repeated) symbols $\rho, \lambda, \beta, \nu$, in some fashion. Is there a subspace $\mathcal{W}$ of $\mathbb{F}^n$ for which this is exactly its “signature”, in the sense described above?

The answer is certainly “not always”. As we have already noted, the last position cannot be a λ -position. So, which signatures are possible?

We begin with an example that provides a partial answer to the question by demonstrating a sufficient condition.

Let

$$\mathcal{W} = \{ (0, a_1, a_2, a_3, a_4, 0, a_1, a_5, a_6, 0, a_4, a_6, a_7, a_7, a_8, a_9, a_8, a_{10}) \mid a_i \in \mathbb{F} \}.$$

Then $\mathcal{W}$ is a subspace of $\mathbb{F}^{18}$ with the following signature:

$$(\nu, \lambda, \beta, \beta, \lambda, \nu, \rho, \beta, \lambda, \nu, \rho, \rho, \lambda, \rho, \lambda, \beta, \rho, \beta).$$

From this example it is quite obvious that the locations of the ν -positions and β -positions are generally unrestricted. Furthermore, if the locations of the λ -positions and ρ -positions correspond to the positions of “matched” left/right parentheses, the whole configuration is achievable as a signature of a subspace of $\mathbb{F}^n$ of the type shown. This describes a sufficient condition, but is it necessary? We claim that it is.

Any configuration of matched left/right parentheses has the property that to the left of any left parenthesis there are more right parentheses than there are left ones. In fact, apart from stating that there are as many left parentheses as there are right parentheses, this property identifies all possible configurations of matched left/right parentheses.

9.2 Theorem All Possible Configurations

A $\rho/\lambda/\beta/\nu$ -configuration of n positions represents a signature of a subspace of $\mathbb{F}^n$ if and only if the number of λ -positions equals the number of ρ -positions, and to the right of every λ -position there are more ρ -positions than λ -positions.

Proof. We focus on the forward implication, since the validity of the reverse implication has already been discussed. Our proof is inductive on n . The claim is trivially true when $n = 1$. If it were not universally true, there would be a smallest $n_o > 1$ for which the claim fails. So, we have some subspace $\mathcal{W}_o$ of $\mathbb{F}^{n_o}$ whose signature does not have the stated property. The signature of the right truncation $\mathcal{U}_o$ of $\mathcal{W}_o$ would enjoy the described property by the minimality of n_o .

This means (by theorem 9.1) that the last position of $\mathcal{W}_o$ would have to be a ρ -position. Yet in that case the signature of $\mathcal{W}_o$ could be recovered from the signature of $\mathcal{U}_o$ by turning a ρ -position into a ν -position or a β -position into a λ -position, and attaching a new ρ -position to the end. Such a transformation could not produce a λ -position that has no more ρ -positions to the right of it than λ -positions, in contradiction to hypothesis about $\mathcal{W}_o$. ■

It is not hard to see that a feasible $\rho/\lambda/\beta/\nu$ -signature does not always describe a unique subspace. How different can two subspaces be if they have identical $\rho/\lambda/\beta/\nu$ -signatures? Note that the persistent zero position may not be the same. Subspaces

$$\mathcal{Z} = \{ (0, a_1, a_2, a_3, a_4, a_4, a_1, a_5, a_6, a_6, a_4, a_6, a_7, a_7, a_8, a_9, a_8, a_{10}) \mid a_i \in \mathbb{F} \}$$

and

$$\mathcal{X} = \{ (0, a_1, a_2, a_3, a_4, a_1, a_1, a_5, a_6, a_4, a_4, a_6, a_7, a_7, a_8, a_9, a_8, a_{10}) \mid a_i \in \mathbb{F} \}$$

of $\mathbb{F}^{18}$ both have the same signature as the subspace $\mathcal{W}$ we had presented earlier.

Of course if a signature begins with a ν , this forces the first position of a subspace to be a zero position. Analogous claim is true when ν is the last symbol of a signature.

One More Thing

So, before we wrap things up, there is one more question that seems too natural to be ignored, and has a simple answer.

If a subspace $\mathcal{Z}$ of $\mathbb{F}^n$, restricted to some k positions, presents as $\mathbb{F}^k$, can one apply a permutation matrix to $\mathcal{Z}$ and have those k positions become **red** positions of the resulting subspace?

The answer is affirmative. Take a permutation matrix that maps those k positions of $\mathbb{F}^n$ to the last k positions (with indices $n-(k-1), \dots, n-1, n$), and apply it to $\mathcal{Z}$ to produce a subspace $\mathcal{Z}_o$. Since, restricted to these last k positions, $\mathcal{Z}_o$ presents as $\mathbb{F}^k$, each of these positions is indeed a **red** position for $\mathcal{Z}_o$.

If k happens to be the dimension of $\mathcal{Z}$, then the last k positions of $\mathcal{Z}_o$ are all of its **red** positions. In this case, the elements of $\mathcal{Z}$ presenting as the standard basis of $\mathbb{F}^k$ on the original k positions, are uniquely determined, and are sent to the **red** basis of $\mathcal{Z}_o$ by the permutation matrix.

10 Concluding Remarks

According to John Baez' insight [2], a good mathematical article should not end abruptly, and we concur.

Originally we were motivated by a search for a palatable argument showing the uniqueness of RREF for the students in the first linear algebra course at our home institution. At the same time we had grown weary of our students fixating on row reduction and the RREF, making these a knee-jerk response to almost any linear-algebraic question. Given the centrality of the orthogonality-related notions in contemporary applications, this just won't do.

Furthermore, it is counterproductive to let students think that proficiency in carrying out Gauss-Jordan elimination or a determinant calculation by hand is a sign of valuable intellectual accomplishment. Present day professionals perform such procedures by hand just about as often as they extract square roots by hand or graph complicated functions with pencils on paper. Proficiency in reducing a quest to Gauss-Jordan elimination (or a determinant calculation), along with an ability to use computing systems, and having the skills to interpret the results, is of course a different matter.

It has to be mentioned that we have only tested "in a jiffy" approach at our home institution, which benefits from having strong and motivated students. It works well and is indeed efficient. The approach is most effective when the emphasis is placed on comprehension of the ideass, and the general proofs are scaled down to more visual and concrete arguments that can be played with in real time.

It is our hope that the right/left-standard bases tack can be adopted more broadly, and that it finds a regular place in the toolbox of modern mathematical educators.

Acknowledgements: We are very grateful to Professors Fernando Gouvea and Scott Taylor, who had kindly read the first version of the article. Their insightful suggestions prompted a significant improvement, as did the blog post by John Baez which Fernando Gouvea had brought to my attention. We would also be remiss not to thank those Colby students whom we had subjected to early versions of the ideas in this article. Their questions and comments contributed greatly to a refinement of the approach.

References

- [1] Grinberg EL. A gauche perspective on row reduced echelon form and its uniqueness. Amer Math Monthly. 2022;129(4):364-73. Available from: <https://doi.org/10.1080/00029890.2022.2027717>.
- [2] Baez J. Why Mathematics Is Boring. 2024. Available from: https://golem.ph.utexas.edu/category/2024/03/why_mathematics_is_boring_1.html.