

Chapter 2

Fields and linear systems

Learning goals

After this chapter you should be able to:

- state the field axioms and check them for $\mathbb{Q}$, $\mathbb{R}$, $\mathbb{C}$ and the finite fields $\mathbb{Z}_p$;
- write a system of linear equations as an augmented matrix and read it in the row picture and the column picture;
- solve a system by Gaussian elimination and back substitution, and bring a matrix to reduced row echelon form;
- describe the solution set of a system, decide existence and uniqueness, and write solutions in parametric form;
- test a set of vectors for linear independence.

2.1 Fields

Lecture 6

Up to now every scalar has been a real number. Looking back at Section 1.2, every rule of vector algebra was proved one entry at a time from a rule of arithmetic for real numbers: commutativity, associativity, the distributive law, and so on. Solving equations needs one more thing: to solve $2x = 6$ we divide by 2, so we must be able to divide by any nonzero number. This section isolates exactly what linear algebra needs from its scalars. A number system with these properties is called a *field*. The real numbers are one field; the rational numbers, the complex numbers and the finite “clock arithmetics” $\mathbb{Z}_p$ are others, and everything in this course that uses only the field rules works in all of them. Finite fields are far from a curiosity: they are the arithmetic behind error-correcting codes (QR codes, data storage) and much of modern cryptography.

2.1.1 The field axioms

Definition 2.1.1 (Field). A *field* is a set $\mathbb{F}$ together with two operations, *addition* $(a, b) \mapsto a + b$ and *multiplication* $(a, b) \mapsto ab$ (also written $a \cdot b$), each taking two elements of $\mathbb{F}$ to an element of $\mathbb{F}$, such that for all $a, b, c \in \mathbb{F}$:

- | | |
|---|---------------------------------|
| (A1) $(a + b) + c = a + (b + c)$ | (addition is associative) |
| (A2) $a + b = b + a$ | (addition is commutative) |
| (A3) there is an element $0 \in \mathbb{F}$ with $a + 0 = a$ for every a | (zero) |
| (A4) for every a there is an element $-a \in \mathbb{F}$ with $a + (-a) = 0$ | (negatives) |
| (M1) $(ab)c = a(bc)$ | (multiplication is associative) |
| (M2) $ab = ba$ | (multiplication is commutative) |
| (M3) there is an element $1 \in \mathbb{F}$ with $1 \neq 0$ and $1a = a$ for every a | (one) |
| (M4) for every $a \neq 0$ there is an element $a^{-1} \in \mathbb{F}$ with $a a^{-1} = 1$ | (inverses) |
| (D) $a(b + c) = ab + ac$ | (distributive law) |

These nine conditions are the *field axioms*. The elements of a field are called *scalars*.

The two lists (A1)–(A4) and (M1)–(M4) say the same things about the two operations, with one difference: every element has a negative, but only nonzero elements are required to have an inverse. Zero never has one (Proposition 2.1.5(a) shows $0 \cdot a = 0 \neq 1$). The condition $1 \neq 0$ rules out the one-element set $\{0\}$. By (M2) and (A2) the laws also hold with the factors in the other order, for example $(b + c)a = ba + ca$ and $0 + a = a$.

Subtraction and division. In a field we define

$$a - b = a + (-b), \quad \frac{a}{b} = a b^{-1} \quad (b \neq 0).$$

So a field is a set in which we can add, subtract, multiply and divide by any nonzero element, with the usual laws of arithmetic.

Example 2.1.2 ($\mathbb{Q}$ and $\mathbb{R}$). The real numbers $\mathbb{R}$ with the usual operations form a field: the nine axioms are the familiar rules of arithmetic. The rational numbers $\mathbb{Q} = \{m/n : m, n \in \mathbb{Z}, n \neq 0\}$ also form a field. The axioms hold in $\mathbb{Q}$ because they hold in $\mathbb{R}$; what needs checking is that the operations never leave $\mathbb{Q}$:

$$\frac{m}{n} + \frac{k}{l} = \frac{ml + kn}{nl}, \quad \frac{m}{n} \cdot \frac{k}{l} = \frac{mk}{nl}, \quad -\frac{m}{n} = \frac{-m}{n}, \quad \left(\frac{m}{n}\right)^{-1} = \frac{n}{m} \quad (m \neq 0),$$

and $0 = 0/1$, $1 = 1/1$ lie in $\mathbb{Q}$. For instance $(2/3)^{-1} = 3/2$.

Example 2.1.3 (Non-examples). The integers $\mathbb{Z}$ satisfy every axiom except (M4): 2 has no inverse in $\mathbb{Z}$, because $2x = 1$ has no integer solution. The natural numbers $\mathbb{N} = \{0, 1, 2, \dots\}$ fail (A4) as well: 1 has no negative in $\mathbb{N}$. Neither is a field. To be a field, a set must be closed under all four operations.

2.1.2 Consequences of the axioms

From now on, a fact about “numbers” is a fact about every field only if it can be proved from the nine axioms. We first show that the special elements in the axioms are unique, so that the notation 0 , 1 , $-a$ and a^{-1} makes sense.

Proposition 2.1.4 (Uniqueness). In a field $\mathbb{F}$:

- (a) there is only one zero and only one one;

(b) each a has only one negative, and each $a \neq 0$ has only one inverse.

Proof. (a) If 0 and $0'$ both satisfy (A3), then $0' = 0' + 0 = 0 + 0' = 0$, using (A3) for 0 , then (A2), then (A3) for $0'$. The same argument with (M3) and (M2) shows $1' = 1$.

(b) Suppose $a + b = 0$ and $a + b' = 0$. Then, using (A3), (A1), (A2) and the two equations,

$$b' = b' + 0 = b' + (a + b) = (b' + a) + b = (a + b') + b = 0 + b = b.$$

For inverses, suppose $ab = 1$ and $ab' = 1$; then $b' = b'1 = b'(ab) = (b'a)b = (ab')b = 1b = b$ by (M3), (M1) and (M2). $\square$

Proposition 2.1.5 (Arithmetic in a field). Let a, b, c be elements of a field $\mathbb{F}$. Then:

- (a) $0a = 0$;
- (b) $(-1)a = -a$;
- (c) $-(-a) = a$, $(-a)b = -(ab)$ and $(-a)(-b) = ab$;
- (d) (no zero divisors) if $ab = 0$, then $a = 0$ or $b = 0$;
- (e) (cancellation) if $a + c = b + c$, then $a = b$; if $ac = bc$ and $c \neq 0$, then $a = b$.

Proof. (a) Since $0 + 0 = 0$ by (A3), the distributive law (D), used with (M2) to put the factor a on the right, gives

$$0a = (0 + 0)a = 0a + 0a.$$

Add $-(0a)$ to both sides. The left side becomes 0 by (A4). On the right, by (A1), (A4) and (A3), $(0a + 0a) + (-(0a)) = 0a + (0a + (-(0a))) = 0a + 0 = 0a$. Hence $0 = 0a$.

(b) By (M3), then (D) with (M2), then (A4) and part (a),

$$a + (-1)a = 1a + (-1)a = (1 + (-1))a = 0a = 0.$$

So $(-1)a$ is a negative of a , and by Proposition 2.1.4(b) it is *the* negative: $(-1)a = -a$.

(c) Since $(-a) + a = a + (-a) = 0$, the element a is a negative of $-a$, so $-(-a) = a$ by uniqueness. Next, $(-a)b = ((-1)a)b = (-1)(ab) = -(ab)$ by (b) and (M1). Finally, $(-a)(-b) = -(a(-b)) = -(-(ab)) = ab$, using the previous identity twice (with (M2)) and $-(-x) = x$.

(d) Suppose $ab = 0$ and $a \neq 0$. Then a^{-1} exists, and by (M1), (M2), (M4), (M3) and part (a),

$$b = 1b = (a^{-1}a)b = a^{-1}(ab) = a^{-1}0 = 0.$$

(e) Add $-c$ to both sides of $a + c = b + c$ and use (A1), (A4), (A3): $a = b$. Similarly, multiply both sides of $ac = bc$ by c^{-1} and use (M1), (M4), (M2) and (M3). $\square$

Part (a) may look too obvious to need a proof, but no axiom says it directly, and it is the reason zero has no inverse. It is also the field version of $0v = 0$ from Section 1.2, proved by the same trick. Part (d) is the property that makes fields so convenient for solving equations:

Proposition 2.1.6 (One equation in one unknown). Let $a, b \in \mathbb{F}$ with $a \neq 0$. The equation $ax = b$ has exactly one solution in $\mathbb{F}$, namely $x = a^{-1}b$.

Proof. $x = a^{-1}b$ is a solution: $a(a^{-1}b) = (aa^{-1})b = 1b = b$. If x_1 and x_2 are both solutions, put $d = x_1 - x_2$. Then $ad = ax_1 + a(-x_2) = ax_1 - ax_2 = b - b = 0$ by (D) and Proposition 2.1.5(c), and since $a \neq 0$, Proposition 2.1.5(d) gives $d = 0$, so $x_1 = x_2$. $\square$

Gaussian elimination (Section 2.3) is this one step, repeated: divide by a nonzero coefficient and subtract. That is why it works over every field. From here on we use the familiar rules of algebra (brackets, signs, fractions, $(a+b)^2 = a^2 + 2ab + b^2$, and so on) in any field, since each of them follows from the axioms by arguments like the ones above.

2.1.3 The complex numbers

Over $\mathbb{R}$ the equation $x^2 = -1$ has no solution, because squares are never negative. The complex numbers add a number i with $i^2 = -1$.

Definition 2.1.7 (Complex numbers). A *complex number* is an expression $z = a + bi$ with $a, b \in \mathbb{R}$; a is its *real part* and b its *imaginary part*. Two complex numbers are equal when their real parts and their imaginary parts agree. The set of complex numbers is $\mathbb{C}$, and its operations are

$$(a + bi) + (c + di) = (a + c) + (b + d)i, \quad (a + bi)(c + di) = (ac - bd) + (ad + bc)i.$$

The *conjugate* of $z = a + bi$ is $\bar{z} = a - bi$.

The multiplication rule is what one gets by expanding the brackets and replacing i^2 by -1 : $(a + bi)(c + di) = ac + adi + bci + bd i^2$. We identify the real number a with $a + 0i$ and write bi for $0 + bi$. Geometrically, $a + bi$ is the point (a, b) of the plane (the *complex plane*), and addition is vector addition (Figure 2.1).

Example 2.1.8. Part by part, $(1 + 2i) + (3 - i) = (1 + 3) + (2 - 1)i = 4 + i$. Expanding,

$$(1 + 2i)(3 - i) = 3 - i + 6i - 2i^2 = (3 + 2) + (-1 + 6)i = 5 + 5i.$$




Figure 2.1: Complex numbers as points of the plane. Adding $1 + 2i$ and $3 - i$ adds their real and imaginary parts, like vectors.

Theorem 2.1.9. $\mathbb{C}$ is a field, with zero $0 = 0 + 0i$, one $1 = 1 + 0i$, negative $-(a + bi) = (-a) + (-b)i$, and, for $a + bi \neq 0$,

$$(a + bi)^{-1} = \frac{a - bi}{a^2 + b^2} = \frac{a}{a^2 + b^2} - \frac{b}{a^2 + b^2}i.$$

Proof. Let $z = a + bi$, $w = c + di$, $u = e + fi$. The addition axioms hold because addition is part by part: (A1) and (A2) follow from the same laws in $\mathbb{R}$ for the real parts and for the imaginary parts, $z + 0 = z$, and $z + (-z) = 0$.

(M2) $zw = (ac - bd) + (ad + bc)i$ and $wz = (ca - db) + (cb + da)i$ agree, by the laws of $\mathbb{R}$.

(M1) Expanding both sides with the definition,

$$\begin{aligned} (zw)u &= ((ac - bd)e - (ad + bc)f) + ((ac - bd)f + (ad + bc)e)i \\ &= (ace - bde - adf - bcf) + (acf - bdf + ade + bce)i, \\ z(wu) &= (a(ce - df) - b(cf + de)) + (a(cf + de) + b(ce - df))i \\ &= (ace - adf - bcf - bde) + (acf + ade + bce - bdf)i, \end{aligned}$$

and the two results have the same real part and the same imaginary part.

(M3) $1z = (1a - 0b) + (1b + 0a)i = z$, and $1 \neq 0$ since their real parts differ.

(D) $z(w+u) = (a(c+e) - b(d+f)) + (a(d+f) + b(c+e))i = ((ac - bd) + (ae - bf)) + ((ad + bc) + (af + be))i = zw + zu$.

(M4) First, $z\bar{z} = (a + bi)(a - bi) = (a^2 + b^2) + (-ab + ba)i = a^2 + b^2$, a real number. If $z \neq 0$, then a and b are not both zero, so $a^2 + b^2 > 0$. Put $z' = (a/(a^2 + b^2)) - (b/(a^2 + b^2))i$; multiplying a complex number by a real number r multiplies both parts by r , so $zz' = z\bar{z}/(a^2 + b^2) = 1$. $\square$

Example 2.1.10. For $z = 1 + 2i$: $\bar{z} = 1 - 2i$ and $z\bar{z} = 1 - 4i^2 = 1 + 4 = 5$. So $(1 + 2i)^{-1} = \frac{1}{5} - \frac{2}{5}i$. Check: $(1 + 2i)(\frac{1}{5} - \frac{2}{5}i) = \frac{1}{5} - \frac{2}{5}i + \frac{2}{5}i - \frac{4}{5}i^2 = \frac{1}{5} + \frac{4}{5} = 1$.

Remark 2.1.11 (No order on $\mathbb{C}$). The real numbers carry an order $<$ that fits the field operations: if $x < y$ then $x + z < y + z$, and if $x, y > 0$ then $xy > 0$. In such an ordered field every nonzero square is positive (if $x < 0$ then $-x > 0$ and $x^2 = (-x)^2 > 0$), so $1 = 1^2 > 0$ and $-1 < 0$. In $\mathbb{C}$, however, $-1 = i^2$ is a nonzero square, so no order on $\mathbb{C}$ fits its operations. Lengths and angles (Section 1.4) rely on the order and the square roots of $\mathbb{R}$, so over $\mathbb{C}$ they need a modification (Chapter 10), and over the finite fields below they have no meaning at all.

Complex numbers return in Chapter 9, where matrices that rotate the plane are best understood over $\mathbb{C}$. Appendix B reviews complex numbers (polar form, $e^{i\theta}$, roots of polynomials) in more detail.

2.1.4 The finite fields $\mathbb{Z}_p$

Arrange the numbers 0, 1, 2, 3, 4 around a clock with five hours (Figure 2.2). Counting round the clock, $3 + 4$ lands on 2: the ordinary sum 7 is replaced by its remainder after division by 5. Multiplying the same way, $3 \cdot 4 = 12$ is replaced by 2.




Figure 2.2: $\mathbb{Z}_5$ as a clock with five hours: four steps on from 3 land on 2, so $3 + 4 = 2$ in $\mathbb{Z}_5$.

For integers x and $n \geq 2$, write $x \bmod n$ for the remainder of x after division by n : the unique $r \in \{0, 1, \dots, n-1\}$ such that $x - r$ is a multiple of n . For example $7 \bmod 5 = 2$ and $-2 \bmod 5 = 3$.

Definition 2.1.12 (Integers modulo n). For an integer $n \geq 2$, the *integers modulo n* are the set $\mathbb{Z}_n = \{0, 1, \dots, n-1\}$ with the operations

$$a \oplus b = (a + b) \bmod n, \quad a \odot b = (ab) \bmod n.$$

We usually write $a + b$ and ab for these operations, and say “in $\mathbb{Z}_n$ ” to make clear which ones are meant.

Proposition 2.1.13. For every $n \geq 2$, $\mathbb{Z}_n$ satisfies all field axioms except possibly (M4), with zero 0, one 1, and negative $-a = n - a$ for $a \neq 0$ (and $-0 = 0$).

Proof. The key fact is: if $x - y$ is a multiple of n , then $x \bmod n = y \bmod n$. Indeed, $x - (x \bmod n)$ is a multiple of n , hence so is $y - (x \bmod n)$, and $x \bmod n$ lies in $\{0, \dots, n-1\}$; by uniqueness of the remainder it equals $y \bmod n$. Since $(x \bmod n) + y$ differs from $x + y$ by a multiple of n , and $(x \bmod n) y$ differs from xy by a multiple of n , it follows that

$$((x \bmod n) + y) \bmod n = (x + y) \bmod n, \quad ((x \bmod n) y) \bmod n = (xy) \bmod n,$$

and likewise with the reduction on y . In words: we may reduce modulo n at any stage of a computation without changing the final remainder. So both sides of each axiom are the remainder of the same integer expression. For example, for (A1), $(a \oplus b) \oplus c = (a + b + c) \bmod n = a \oplus (b \oplus c)$, and for (D), $a \odot (b \oplus c) = (a(b + c)) \bmod n = (ab + ac) \bmod n = (a \odot b) \oplus (a \odot c)$. Axioms (A2), (M1), (M2) follow in the same way from the integer laws. For (A3) and (M3), $a + 0 = a$ and $1 \cdot a = a$ already lie in $\{0, \dots, n-1\}$, and $1 \neq 0$ because $n \geq 2$. For (A4), $a + (n - a) = n$, and $n \bmod n = 0$. $\square$

So the only question is whether nonzero elements have inverses. To find the inverse of a , look along the row of a in the multiplication table for an entry 1.

Example 2.1.14 ($\mathbb{Z}_5$ and $\mathbb{Z}_6$). The multiplication tables of $\mathbb{Z}_5$ and $\mathbb{Z}_6$ are

$\cdot$	0	1	2	3	4
0	0	0	0	0	0
1	0	1	2	3	4
2	0	2	4	1	3
3	0	3	1	4	2
4	0	4	3	2	1

$\cdot$	0	1	2	3	4	5
0	0	0	0	0	0	0
1	0	1	2	3	4	5
2	0	2	4	0	2	4
3	0	3	0	3	0	3
4	0	4	2	0	4	2
5	0	5	4	3	2	1

In $\mathbb{Z}_5$ every nonzero row contains a 1: $1^{-1} = 1$, $2^{-1} = 3$ (as $2 \cdot 3 = 6$ and $6 \bmod 5 = 1$), $3^{-1} = 2$ and $4^{-1} = 4$ (as $4 \cdot 4 = 16$ and $16 \bmod 5 = 1$). Each nonzero row is a rearrangement of 0, 1, 2, 3, 4. So $\mathbb{Z}_5$ is a field. In $\mathbb{Z}_6$ we have $2 \cdot 3 = 0$ with $2, 3 \neq 0$, which is impossible in a field by Proposition 2.1.5(d). Accordingly the row of 2 reads 0, 2, 4, 0, 2, 4 and contains no 1: 2 has no inverse, and $\mathbb{Z}_6$ is not a field.

The difference between 5 and 6 is that 5 is prime. To prove that this is the whole story we need one fact about primes.

Proposition 2.1.15 (Euclid’s lemma). Let p be a prime number and a, b integers. If p divides ab , then p divides a or p divides b .

Proof. Suppose $p \mid ab$ and $p \nmid a$; we show $p \mid b$. Let S be the set of positive integers x such that $p \mid xb$. Both p and $|a|$ lie in S ($|a| \neq 0$ since $p \nmid a$). Let d be the smallest element of S . Every $x \in S$ is a multiple of d : write $x = qd + r$ with $0 \leq r < d$; then $rb = xb - qdb$ is divisible by p , so if $r > 0$ then $r \in S$, contradicting the minimality of d . Hence $r = 0$. In particular d divides p , so $d = 1$ or $d = p$. If $d = p$, then p divides $|a|$, which is excluded. So $d = 1$, and $1 \in S$ says that $p \mid b$. $\square$

Theorem 2.1.16. $\mathbb{Z}_n$ is a field if and only if n is prime.

Proof. Suppose n is not prime. Then $n = ab$ with integers $2 \leq a, b \leq n-1$. These are nonzero elements of $\mathbb{Z}_n$, but $a \odot b = n \bmod n = 0$. A field has no zero divisors (Proposition 2.1.5(d)), so $\mathbb{Z}_n$ is not a field.

Now let $n = p$ be prime. By Proposition 2.1.13 it remains to prove (M4). Take $a \in \mathbb{Z}_p$ with $a \neq 0$, and consider the p products $a \odot x$ for $x = 0, 1, \dots, p-1$. They are all different: if $a \odot x = a \odot y$ with $x, y \in \mathbb{Z}_p$, then $ax - ay = a(x - y)$ is a multiple of p . By Euclid's lemma p divides a or $x - y$. It cannot divide a , since $1 \leq a \leq p-1$; so p divides $x - y$, and since $-(p-1) \leq x - y \leq p-1$, this forces $x = y$. Hence the p products are p different elements of the p -element set $\mathbb{Z}_p$, so every element appears exactly once; in particular $a \odot x = 1$ for some x , and that x is a^{-1} . $\square$

Example 2.1.17 ($\mathbb{Z}_2$). The smallest field is $\mathbb{Z}_2 = \{0, 1\}$, with

$+$	0	1
0	0	1
1	1	0

$\cdot$	0	1
0	0	0
1	0	1

Here $1 + 1 = 0$, so $-1 = 1$ and subtraction is the same as addition. Reading 0 as “even” and 1 as “odd”, the tables are the rules for adding and multiplying even and odd numbers. Reading them as bits, addition is the *exclusive or* (XOR) and multiplication is *and*. Linear algebra over $\mathbb{Z}_2$ is the basis of the error-correcting codes used to store and transmit data.

2.1.5 Linear algebra over a field

Definition 2.1.18 (The space $\mathbb{F}^n$). Let $\mathbb{F}$ be a field. $\mathbb{F}^n$ is the set of columns $\mathbf{v} = (v_1, \dots, v_n)$ with entries $v_i \in \mathbb{F}$. Addition and scalar multiplication are entry by entry, as in $\mathbb{R}^n$: $(\mathbf{v} + \mathbf{w})_i = v_i + w_i$ and $(c\mathbf{v})_i = c v_i$ for $c \in \mathbb{F}$.

The eight rules of vector algebra (Theorem 1.2.7) hold in $\mathbb{F}^n$ for every field $\mathbb{F}$, with the same proof: each rule was checked entry by entry using one law of arithmetic, and each of those laws is a field axiom or, for rule (8), the distributive law (D) combined with (M2) (see the remark after Definition 2.1.1). The same goes for linear combinations and span (Section 1.3), whose definitions use only these operations.

Example 2.1.19. In $\mathbb{Z}_5^2$, coordinate by coordinate, $(3, 4) + (4, 2) = (7 \bmod 5, 6 \bmod 5) = (2, 1)$.

Now solve the system over $\mathbb{Z}_5$

$$\begin{aligned} x + y &= 1 \\ x + 2y &= 4 \end{aligned}$$

Subtracting the first equation from the second gives $y = 4 - 1 = 3$. Then $x = 1 - y = 1 - 3 = -2$, and $-2 = 3$ in $\mathbb{Z}_5$. So $(x, y) = (3, 3)$. Check modulo 5: $3 + 3 = 6$, which is 1; and $3 + 2 \cdot 3 = 9$, which is 4.

Finally, $2x = 3$ in $\mathbb{Z}_5$ has the unique solution $x = 2^{-1} \cdot 3 = 3 \cdot 3 = 9 \bmod 5 = 4$ (Proposition 2.1.6); indeed $2 \cdot 4 = 8$, which is 3.

The methods of the next two sections, writing a system as an augmented matrix and solving it by elimination, use only the four field operations, so they work over every field. Lengths, angles and the dot-product geometry of Section 1.4 are the exception (Remark 2.1.11). In Chapter 4 vector spaces are defined over an arbitrary field $\mathbb{F}$; until then, “scalar” still means a real number unless a field is named.

Check yourself

- (a) Find 3^{-1} in $\mathbb{Z}_7$ and solve $3x + 2 = 6$ in $\mathbb{Z}_7$. (b) Compute $(2 + i)^{-1}$ in $\mathbb{C}$ and check your answer. (c) Is $\mathbb{Z}_9$ a field?

Supplement: beyond the lecture

This material is not needed later in the course, but it rounds out the picture.

Inverses by Fermat’s little theorem. In $\mathbb{Z}_p$ the nonzero elements are $1, \dots, p-1$. Multiplying each of them by a fixed $a \neq 0$ only rearranges them (proof of Theorem 2.1.16), so the product of all of them is unchanged: $a^{p-1} (1 \cdot 2 \cdots (p-1)) = 1 \cdot 2 \cdots (p-1)$ in $\mathbb{Z}_p$. Cancelling the nonzero product (Proposition 2.1.5(e)) gives $a^{p-1} = 1$, so $a^{-1} = a^{p-2}$. For example, in $\mathbb{Z}_5$, $2^{-1} = 2^3 = 8 \bmod 5 = 3$.

The characteristic. In $\mathbb{Z}_p$, adding 1 to itself p times gives 0; in $\mathbb{Q}$, $\mathbb{R}$ and $\mathbb{C}$ no such sum is ever 0. The smallest number of ones that add up to 0 (if there is one) is the *characteristic* of the field, and it is always a prime: if $(ab) \cdot 1 = (a \cdot 1)(b \cdot 1) = 0$ with a, b smaller, one of the factors would already be 0. Over a field of characteristic 2, $1 + 1 = 0$, so one cannot divide by 2, and a formula such as the midpoint $\frac{1}{2}(a + b)$ has no meaning there.

Other finite fields. $\mathbb{Z}_4$ is not a field ($2 \cdot 2 = 0$), but there is a field with four elements, built differently: $\mathbb{F}_4 = \{0, 1, \omega, \omega + 1\}$ with $1 + 1 = 0$ and $\omega^2 = \omega + 1$. For example $\omega(\omega + 1) = \omega^2 + \omega = (\omega + 1) + \omega = 1$, so ω and $\omega + 1$ are inverse to each other. In general there is a field with q elements exactly when q is a power of a prime, and it is essentially unique. These fields are the arithmetic of Reed–Solomon codes, used in QR codes and in data storage.

Fields between $\mathbb{Q}$ and $\mathbb{R}$. The set $\mathbb{Q}(\sqrt{2}) = \{a + b\sqrt{2} : a, b \in \mathbb{Q}\}$ is a field: it is closed under the operations, and $(a + b\sqrt{2})^{-1} = (a - b\sqrt{2})/(a^2 - 2b^2)$, where $a^2 - 2b^2 \neq 0$ for $(a, b) \neq (0, 0)$ because $\sqrt{2}$ is irrational. The same conjugate trick inverts complex numbers.

2.2 Systems of linear equations

Lecture 7

Twice already a question about vectors has turned into a system of linear equations. In Section 1.1 we asked which combination of two columns reaches a target, and in Section 1.3 we asked whether a vector lies in a span. In both cases the question was really about numbers that satisfy several linear equations at once. This section makes the notion precise, looks at a system in two complementary ways, the *row picture* and the *column picture*, proves that a system has no solution, exactly one, or infinitely many, and introduces the *augmented matrix*, the compact record of a system that the rest of the chapter works with. Throughout, $\mathbb{F}$ is a field (Section 2.1); in the pictures it is $\mathbb{R}$.

2.2.1 Linear equations

Definition 2.2.1 (Linear equation). A *linear equation* in the unknowns $x_1, \dots, x_n$ is an equation of the form

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = b,$$

where the *coefficients* $a_1, \dots, a_n$ and the *constant term* b are elements of $\mathbb{F}$. A *solution* of the equation is a list $(s_1, \dots, s_n) \in \mathbb{F}^n$ such that $a_1s_1 + \cdots + a_ns_n = b$.

Each unknown appears to the first power only, multiplied by a constant, and the products are added. There are no squares or higher powers of unknowns, no products of two unknowns, and no unknowns inside other functions.

Example 2.2.2. The equation $2x + 3y - z = 5$ is linear, with coefficients 2, 3, -1 and constant term 5. The list $(1, 1, 0)$ is a solution, since $2 \cdot 1 + 3 \cdot 1 - 0 = 5$; the list $(0, 0, 0)$ is not. The equation $3x = y + 1$ is linear too: moving y to the left gives the standard form $3x - y = 1$. The equations $x^2 + y = 1$, $xy = 4$, $\sin x + y = 0$ and $\sqrt{x} = y$ are not linear.

Over $\mathbb{R}$, the solutions of one linear equation in two unknowns (with not both coefficients zero) form a line in the plane, and in three unknowns they form a plane in space whose normal vector is the vector of coefficients (Section 1.5): the equation $a_1x_1 + a_2x_2 + a_3x_3 = b$ says $\mathbf{n} \cdot \mathbf{x} = b$ with $\mathbf{n} = (a_1, a_2, a_3)$. If all coefficients are zero, the left side is 0 for every choice of the unknowns. Then the equation $0x_1 + \cdots + 0x_n = b$ has no solution when $b \neq 0$, and every list in $\mathbb{F}^n$ is a solution when $b = 0$.

2.2.2 Systems and their solutions

Definition 2.2.3 (System of linear equations). A *system of m linear equations in n unknowns* is a list of linear equations in the same unknowns $x_1, \dots, x_n$:

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1, \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2, \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m. \end{aligned}$$

The coefficient a_{ij} stands in equation i in front of unknown j : the first index names the equation (the row), the second the unknown (the column). A *solution* of the system is a list $(s_1, \dots, s_n) \in \mathbb{F}^n$ that is a solution of every equation at once. The set of all solutions is the *solution set*. The system is *consistent* if it has at least one solution and *inconsistent* if it has none.

Example 2.2.4. Consider the system

$$\begin{aligned} x + y + z &= 6, \\ 2x - y + z &= 3, \\ x + 2y - z &= 2. \end{aligned} \tag{2.1}$$

The list $(1, 2, 3)$ is a solution: substituting $x = 1$, $y = 2$, $z = 3$ gives $1 + 2 + 3 = 6$, then $2 \cdot 1 - 2 + 3 = 3$, then $1 + 2 \cdot 2 - 3 = 2$, and all three equations hold. The list $(6, 0, 0)$ satisfies the first equation, since $6 + 0 + 0 = 6$, but not the second, since $2 \cdot 6 - 0 + 0 = 12 \neq 3$; so it is not a solution. Satisfying some of the equations is not enough. (Gaussian elimination, Section 2.3, shows that $(1, 2, 3)$ is the only solution.)

Definition 2.2.5 (Equivalent systems). Two systems of linear equations in the same unknowns are *equivalent* if they have the same solution set.

Solving a system means replacing it, step by step, by simpler equivalent systems until the solutions can be read off, for example the system $x = 1, y = 2, z = 3$. Section 2.3 (Gaussian elimination) supplies the steps.

2.2.3 The row picture

In the *row picture* we read a system one equation at a time. Each equation has its own solution set, and the solution set of the system is the intersection of these sets: the points that lie in all of them.

Example 2.2.6. Over $\mathbb{R}$, take the system $2x + y = 5, x - y = 1$. Each equation describes a line in the plane, and the lines cross at the single point $(2, 1)$; indeed $2 \cdot 2 + 1 = 5$ and $2 - 1 = 1$ (Figure 2.3, left).

If the second equation is changed to $2x + y = 1$, the two lines have the same slope -2 and are parallel, so the system has no solution. Algebraically the two left sides are equal, so a solution would force $5 = 1$ (Figure 2.3, middle).

If the second equation is changed to $4x + 2y = 10$, it is twice the first equation and describes the same line. Every point of that line is a solution, and there are infinitely many (Figure 2.3, right).




Figure 2.3: The row picture of Example 2.2.6: two lines in the plane meet in one point, in no point, or in infinitely many points.

With three unknowns over $\mathbb{R}$, each equation (with not all coefficients zero) describes a plane in space. The three planes of system (2.1) all pass through $(1, 2, 3)$, and Section 2.3 shows that they meet only there. Three planes can also share a whole line, like the pages of an open book, or have no common point at all even though no two of them are parallel: each pair then meets in a line, and the three lines are parallel, like the three edges of a tent (Figure 2.4). Two parallel planes, of course, already have no common point.

2.2.4 No solution, one, or infinitely many

In all these pictures the number of solutions was 0, 1 or infinite. This is a theorem. To prove it we write each equation with a dot product: if $\mathbf{r}_i = (a_{i1}, \dots, a_{in})$ is the *row* of coefficients of equation i , then its left side is $a_{i1}x_1 + \dots + a_{in}x_n = \mathbf{r}_i \cdot \mathbf{x}$, where $\mathbf{x} = (x_1, \dots, x_n)$, so equation i reads $\mathbf{r}_i \cdot \mathbf{x} = b_i$. The dot




Figure 2.4: Two ways three planes can meet other than in a single point.

product on $\mathbb{F}^n$ is defined by the same formula as on $\mathbb{R}^n$, and parts (a) to (c) of Theorem 1.4.6 hold over any field with the same proofs, since they use only the field axioms.

Theorem 2.2.7 (Zero, one or infinitely many). Let $\mathbb{F}$ be $\mathbb{R}$, or more generally any field with infinitely many elements, such as $\mathbb{Q}$ or $\mathbb{C}$. A system of linear equations over $\mathbb{F}$ has either no solution, exactly one solution, or infinitely many solutions. More precisely, if $\mathbf{s} \neq \mathbf{t}$ are two solutions, then $\mathbf{s} + c(\mathbf{t} - \mathbf{s})$ is a solution for every $c \in \mathbb{F}$, and these solutions are distinct for distinct c .

Proof. It suffices to prove the second statement: a system with at least two solutions then has at least as many solutions as $\mathbb{F}$ has elements, hence infinitely many.

Let $\mathbf{s}$ and $\mathbf{t}$ be solutions with $\mathbf{s} \neq \mathbf{t}$, so that $\mathbf{r}_i \cdot \mathbf{s} = b_i$ and $\mathbf{r}_i \cdot \mathbf{t} = b_i$ for every $i = 1, \dots, m$. Fix $c \in \mathbb{F}$ and put $\mathbf{p} = \mathbf{s} + c(\mathbf{t} - \mathbf{s})$. By additivity and homogeneity of the dot product,

$$\mathbf{r}_i \cdot \mathbf{p} = \mathbf{r}_i \cdot \mathbf{s} + c(\mathbf{r}_i \cdot \mathbf{t} - \mathbf{r}_i \cdot \mathbf{s}) = b_i + c(b_i - b_i) = b_i + c \cdot 0 = b_i.$$

So $\mathbf{p}$ solves equation i , for every i , and is a solution of the system.

If $c \neq c'$ give the same point, then $c(\mathbf{t} - \mathbf{s}) = c'(\mathbf{t} - \mathbf{s})$, so $(c - c')(\mathbf{t} - \mathbf{s}) = \mathbf{0}$. Since $c - c' \neq 0$, multiplying by $(c - c')^{-1}$ gives $\mathbf{t} - \mathbf{s} = \mathbf{0}$, contradicting $\mathbf{s} \neq \mathbf{t}$. Hence distinct values of c give distinct solutions. $\square$

Over $\mathbb{R}$ the solutions $\mathbf{s} + c(\mathbf{t} - \mathbf{s})$ fill the whole line through $\mathbf{s}$ and $\mathbf{t}$: as soon as a system has two solutions, it has a line of them. In particular no system has exactly two, or exactly seven, solutions.

Remark 2.2.8. The hypothesis that $\mathbb{F}$ is infinite is needed. Over the field $\mathbb{Z}_2 = \{0, 1\}$ of Section 2.1, the equation $x + y = 1$ has exactly two solutions among the four candidates $(0, 0)$, $(1, 0)$, $(0, 1)$, $(1, 1)$, namely $(1, 0)$ and $(0, 1)$; the candidate $(1, 1)$ fails because $1 + 1 = 0$ in $\mathbb{Z}_2$. The proof above still shows that two solutions $\mathbf{s} \neq \mathbf{t}$ produce the solutions $\mathbf{s} + c(\mathbf{t} - \mathbf{s})$, but over $\mathbb{Z}_p$ there are only p values of c .

2.2.5 The column picture

In the *column picture* we read the system one unknown at a time. Collect the coefficients of unknown j into the column $\mathbf{a}_j = (a_{1j}, \dots, a_{mj}) \in \mathbb{F}^m$, and the constant terms into the right-hand side $\mathbf{b} = (b_1, \dots, b_m) \in \mathbb{F}^m$. Then the system of Definition 2.2.3 is the single vector equation

$$x_1 \mathbf{a}_1 + x_2 \mathbf{a}_2 + \dots + x_n \mathbf{a}_n = \mathbf{b}. \quad (2.2)$$

Indeed, entry i of the left side is $a_{i1}x_1 + \dots + a_{in}x_n$, and two vectors are equal exactly when their entries are equal, so (2.2) says precisely that equation i holds for every i .

Proposition 2.2.9. A list $(s_1, \dots, s_n)$ solves the system exactly when $s_1 \mathbf{a}_1 + \dots + s_n \mathbf{a}_n = \mathbf{b}$. Consequently the system is consistent if and only if

$$\mathbf{b} \in \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_n\}.$$

Proof. The first statement is the entry-by-entry comparison just made. For the second: the system has a solution exactly when some weights $s_1, \dots, s_n$ combine the columns into $\mathbf{b}$, that is, when $\mathbf{b}$ is a linear combination of $\mathbf{a}_1, \dots, \mathbf{a}_n$, which is the definition of $\mathbf{b} \in \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ (Section 1.3). $\square$

Example 2.2.10. For $2x + y = 5$, $x - y = 1$ the columns are $\mathbf{a}_1 = (2, 1)$ and $\mathbf{a}_2 = (1, -1)$, and $\mathbf{b} = (5, 1)$:

$$x \begin{bmatrix} 2 \\ 1 \end{bmatrix} + y \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \begin{bmatrix} 5 \\ 1 \end{bmatrix}.$$

The weights $x = 2$, $y = 1$ work: $2(2, 1) = (4, 2)$, and coordinate by coordinate $4 + 1 = 5$ and $2 + (-1) = 1$, so $2\mathbf{a}_1 + \mathbf{a}_2 = (5, 1) = \mathbf{b}$ (Figure 2.5). The columns do not lie on one line through the origin, so by Theorem 1.3.8 their span is all of $\mathbb{R}^2$, and every right-hand side is reached. It is reached in only one way: if also $x'\mathbf{a}_1 + y'\mathbf{a}_2 = \mathbf{b}$, subtracting gives $(x - x')\mathbf{a}_1 + (y - y')\mathbf{a}_2 = \mathbf{0}$, that is $2(x - x') + (y - y') = 0$ and $(x - x') - (y - y') = 0$; adding the two equations gives $3(x - x') = 0$, so $x = x'$, and then $y = y'$.

In the parallel case $2x + y = 5$, $2x + y = 1$ the columns are $(2, 2)$ and $(1, 1)$; both lie on the line spanned by $(1, 1)$, and $\mathbf{b} = (5, 1)$ does not (its entries differ), so the system is inconsistent. In the case of the repeated line $2x + y = 5$, $4x + 2y = 10$ the columns are $(2, 4)$ and $(1, 2)$ and $\mathbf{b} = (5, 10)$ lies on their line; for example $2(2, 4) + 1(1, 2) = (5, 10)$ and $0(2, 4) + 5(1, 2) = (5, 10)$, and there are infinitely many such pairs of weights.

For system (2.1), the solution $(1, 2, 3)$ gives the weights of a combination of three columns in $\mathbb{R}^3$:

$$1 \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} + 2 \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix} + 3 \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix} = \begin{bmatrix} 1 + 2 + 3 \\ 2 - 2 + 3 \\ 1 + 4 - 3 \end{bmatrix} = \begin{bmatrix} 6 \\ 3 \\ 2 \end{bmatrix},$$

and the three entries are exactly the three checks of Example 2.2.4.




Figure 2.5: The column picture of Example 2.2.10: two copies of $\mathbf{a}_1$ and one of $\mathbf{a}_2$ reach $\mathbf{b}$.

The two pictures live in different spaces. In the row picture a point is a choice of the n unknowns, so it lies in $\mathbb{R}^n$. In the column picture the columns and $\mathbf{b}$ have one entry per equation, so they lie in $\mathbb{R}^m$. The row picture is the easier one to see for two or three unknowns; the column picture links systems to spans, and it leads to linear independence (Section 2.6) and to the product $A\mathbf{x}$ (Section 3.1).

2.2.6 The augmented matrix

When we solve a system, the names of the unknowns play no role; only the coefficients and the constant terms matter.

Definition 2.2.11 (Matrix, coefficient matrix, augmented matrix). An $m \times n$ matrix over $\mathbb{F}$ is a rectangular array of elements of $\mathbb{F}$ with m rows and n columns. The entry in row i and column j is written a_{ij} . The *coefficient matrix* of the system in Definition 2.2.3 is the $m \times n$ matrix $A = (a_{ij})$, with one row per equation and one column per unknown; its *augmented matrix* is the $m \times (n+1)$ matrix obtained by adding the column $\mathbf{b}$ of constant terms,

$$A = \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{m1} & \cdots & a_{mn} \end{bmatrix}, \quad [A \mid \mathbf{b}] = \left[\begin{array}{ccc|c} a_{11} & \cdots & a_{1n} & b_1 \\ \vdots & & \vdots & \vdots \\ a_{m1} & \cdots & a_{mn} & b_m \end{array} \right].$$

The vertical bar only marks where the constant terms begin.

Row i of $[A \mid \mathbf{b}]$ records equation i , the object of the row picture; column $j \leq n$ is the column $\mathbf{a}_j$ of unknown x_j , the object of the column picture; and the last column is $\mathbf{b}$. Two points need care when a system is written down: an unknown missing from an equation has coefficient 0, which must be written; and the unknowns must be taken in the same order in every equation.

Example 2.2.12. System (2.1) has augmented matrix

$$\left[\begin{array}{ccc|c} 1 & 1 & 1 & 6 \\ 2 & -1 & 1 & 3 \\ 1 & 2 & -1 & 2 \end{array} \right].$$

The system $2x - z = 1$, $3y + x = 7$, $z = 4 - y$ is first put in standard form, unknowns in the order x, y, z on the left and constants on the right: $2x + 0y - z = 1$, $x + 3y + 0z = 7$, $0x + y + z = 4$. Its augmented matrix is

$$\left[\begin{array}{ccc|c} 2 & 0 & -1 & 1 \\ 1 & 3 & 0 & 7 \\ 0 & 1 & 1 & 4 \end{array} \right].$$

Conversely, a matrix can be read back as a system. A row $[0 \ 0 \ 0 \mid 5]$ is the equation $0x + 0y + 0z = 5$, that is $0 = 5$, so any system with such a row is inconsistent. More generally, a row $[0 \ \cdots \ 0 \mid c]$ with $c \neq 0$ makes a system inconsistent.

In Section 2.3 we solve systems by working only on the rows of the augmented matrix, with three row operations that never change the solution set; in Section 3.1 the coefficient matrix acts on the vector of unknowns and the whole system becomes the single equation $A\mathbf{x} = \mathbf{b}$.

Check yourself

For the system $x + 2y = 4$, $3x - y = 5$: (a) write its augmented matrix; (b) check that $(2, 1)$ is a solution; (c) write the system as a vector equation in the column picture and verify the solution there.

Supplement: beyond the lecture

This material is not needed later in the course, but it rounds out the picture.

Counting solutions over a finite field. Over $\mathbb{Z}_p$ a system cannot have infinitely many solutions, since $\mathbb{Z}_p^n$ has only p^n elements. Theorem 2.2.7 has a finite-field version: a consistent system over $\mathbb{Z}_p$ has exactly p^k solutions, where $k \geq 0$ is the number of free variables, introduced in Section 2.4. For example, $x + y + z = 1$ over $\mathbb{Z}_2$ has one free choice for each of y and z , and so $2^2 = 4$ solutions: $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ and $(1, 1, 1)$.

Why “linear”. If $\mathbf{s}$ solves the system with right-hand side $\mathbf{b}$ and $\mathbf{s}'$ solves the system with the same coefficients and right-hand side $\mathbf{b}'$, then $\mathbf{s} + \mathbf{s}'$ solves it with right-hand side $\mathbf{b} + \mathbf{b}'$, and $c\mathbf{s}$ solves it with right-hand side $c\mathbf{b}$. This follows from the same dot-product rules used in the proof of Theorem 2.2.7. Responses add and scale with their causes; this *superposition principle* is the defining feature of linear problems in physics and engineering, and it fails as soon as a square or a product of unknowns appears.

A very old method. Systems of linear equations, and the idea of recording only their numbers in a rectangular array, are ancient. A Chinese mathematical text compiled about two thousand years ago, *The Nine Chapters on the Mathematical Art*, solves systems by arranging the coefficients in columns of counting rods and combining the columns, which is the elimination method of Section 2.3 in all but name.

2.3 Gaussian elimination

Lecture 8

In the previous section we wrote a system of linear equations compactly as an augmented matrix $[A \mid \mathbf{b}]$. This section gives a method that solves *every* such system, in the same order every time, and that a computer can carry out: *Gaussian elimination*. It has two stages. *Forward elimination* uses three simple moves, the elementary row operations, to bring the matrix into a staircase shape called row echelon form. *Back substitution* then solves the staircase system from the bottom up, one unknown at a time. The theorem that makes the method trustworthy is that row operations never change the set of solutions.

The idea is the one you know from school. To solve

$$x + y = 3, \quad x - y = 1,$$

subtract the first equation from the second: x disappears and $-2y = -2$, so $y = 1$. Substituting into the first equation gives $x + 1 = 3$, so $x = 2$. Gaussian elimination organises this trick so that it works for any number of equations and unknowns. The method is named after Carl Friedrich Gauss, but it is about two thousand years older: the Chinese text *Nine Chapters on the Mathematical Art* solves systems in exactly this way, with the coefficients laid out in columns on a counting board.

2.3.1 Elementary row operations

Each row of an augmented matrix is one equation: its coefficients, then, after the bar, its right-hand side. We write $R_1, R_2, \dots, R_m$ for the rows.

Definition 2.3.1 (Elementary row operations). The *elementary row operations* on a matrix are:

- (1) *swap*: exchange two rows, written $R_i \leftrightarrow R_j$;
- (2) *scaling*: multiply a row by a *nonzero* scalar c , written $R_i \rightarrow c R_i$;

(3) *replacement*: add c times one row to a *different* row, written $R_i \rightarrow R_i + c R_j$ ($i \neq j$).

The arrow $\rightarrow$ is read “is replaced by”.

Each operation acts on whole rows, including the entries after the bar: changing the left side of an equation without its right side would change the equation.

Example 2.3.2. Start from the system above. A swap lists the same two equations in the other order:

$$\left[\begin{array}{cc|c} 1 & 1 & 3 \\ 1 & -1 & 1 \end{array} \right] \xrightarrow{R_1 \leftrightarrow R_2} \left[\begin{array}{cc|c} 1 & -1 & 1 \\ 1 & 1 & 3 \end{array} \right].$$

Scaling the second row by 2 multiplies each entry by 2: $2 \cdot 1 = 2$, $2 \cdot (-1) = -2$, $2 \cdot 1 = 2$, giving the equation $2x - 2y = 2$, which has the same solutions as $x - y = 1$. The replacement $R_2 \rightarrow R_2 - R_1$ is the school trick:

$$\left[\begin{array}{cc|c} 1 & 1 & 3 \\ 1 & -1 & 1 \end{array} \right] \xrightarrow{R_2 \rightarrow R_2 - R_1} \left[\begin{array}{cc|c} 1 & 1 & 3 \\ 0 & -2 & -2 \end{array} \right],$$

since $1 - 1 = 0$, $-1 - 1 = -2$ and $1 - 3 = -2$. The new second row says $-2y = -2$: the unknown x has been eliminated from it.

2.3.2 Row operations keep the solutions

Definition 2.3.3 (Row equivalence). Two matrices of the same size are *row equivalent* if one can be obtained from the other by a finite sequence of elementary row operations.

Theorem 2.3.4 (Row operations keep the solution set). If two augmented matrices are row equivalent, then the corresponding linear systems have exactly the same solutions.

Proof. A sequence of row operations is one operation after another, so it is enough to prove the claim for a single elementary row operation. Write the i -th equation of the old system as $L_i(\mathbf{x}) = b_i$, where $L_i(\mathbf{x}) = a_{i1}x_1 + \cdots + a_{in}x_n$ is its left side.

No solutions are lost. Let $\mathbf{x}$ be a solution of the old system. A swap only reorders the equations, so $\mathbf{x}$ still satisfies all of them. A scaling replaces $L_i(\mathbf{x}) = b_i$ by $c L_i(\mathbf{x}) = c b_i$, which $\mathbf{x}$ satisfies. For a replacement $R_i \rightarrow R_i + c R_j$, the new i -th row has coefficients $a_{ik} + c a_{jk}$ and right-hand side $b_i + c b_j$. Since $L_i(\mathbf{x}) = b_i$ and $L_j(\mathbf{x}) = b_j$,

$$(a_{i1} + c a_{j1})x_1 + \cdots + (a_{in} + c a_{jn})x_n = L_i(\mathbf{x}) + c L_j(\mathbf{x}) = b_i + c b_j,$$

so $\mathbf{x}$ satisfies the new row i ; all other rows are unchanged. Hence every solution of the old system is a solution of the new one.

No solutions are gained. Every elementary row operation can be undone by an elementary row operation of the same type: $R_i \leftrightarrow R_j$ is undone by $R_i \leftrightarrow R_j$; $R_i \rightarrow c R_i$ is undone by $R_i \rightarrow \frac{1}{c} R_i$, which is allowed because $c \neq 0$; and $R_i \rightarrow R_i + c R_j$ is undone by $R_i \rightarrow R_i - c R_j$ (row j was not changed, since $j \neq i$). So the old system is obtained from the new one by an elementary row operation, and by the first part every solution of the new system is a solution of the old one.

The two solution sets are therefore equal. $\square$

The two conditions in Definition 2.3.1 are exactly what the proof needs. Multiplying a row by $c = 0$ would turn its equation into $0 = 0$; that information could never be recovered, and the new system could have more solutions (for $x + y = 3$, $x - y = 1$, scaling the second row by 0 leaves only $x + y = 3$, which has infinitely many solutions). Similarly, “ $R_i \rightarrow R_i + c R_i$ ” is excluded because it is the scaling by $1 + c$, which is not allowed when $c = -1$.

This theorem looks obvious, but it is the reason that every answer found by elimination is an answer to the original problem. It is used every time a matrix is row reduced, for the rest of the course.

2.3.3 Row echelon form

Definition 2.3.5 (Row echelon form). In a nonzero row of a matrix, the first entry that is not zero is the *leading entry* of the row. A matrix is in *row echelon form* if

- (i) all rows consisting entirely of zeros are below all nonzero rows, and
- (ii) the leading entry of each nonzero row lies in a column strictly to the right of the leading entry of the row above it.

The leading entries of a matrix in row echelon form are called its *pivots*, and the columns that contain them are its *pivot columns*.

Condition (ii) implies that every entry below a pivot is zero: the pivots step down and to the right like a staircase, and everything under the staircase is zero (Figure 2.6). Pivots are studied in detail in Section 2.4.

■	*	*	*	*	*	*
0	0	■	*	*	*	*
0	0	0	■	*	*	*
0	0	0	0	0	0	0

■ = pivot (nonzero)
* = any number

Figure 2.6: The shape of a row echelon form of a 4×7 augmented matrix: the pivots step to the right, everything under the staircase is zero, and the zero row is at the bottom.

Example 2.3.6.

$$\begin{bmatrix} 2 & 1 & 0 & 3 \\ 0 & 0 & 5 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 1 & 4 & 2 & 0 \\ 0 & 3 & 1 & 2 \\ 0 & 6 & 0 & 1 \end{bmatrix} \quad \begin{bmatrix} 1 & 2 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 4 & 3 \end{bmatrix}$$

The first matrix is in row echelon form, with pivots 2 and 5 in columns 1 and 3. The second is not: the leading entries of rows 2 and 3 are both in column 2. The third is not: its zero row lies above a nonzero row.

A system whose augmented matrix is in row echelon form can be solved from the bottom up, one unknown at a time. Getting there is the job of forward elimination.

2.3.4 Forward elimination

Theorem 2.3.7 (Existence of a row echelon form). Every matrix is row equivalent to a matrix in row echelon form.

The proof is the algorithm itself.

Forward elimination. Given an $m \times n$ matrix:

1. Find the leftmost column that is not entirely zero. If there is none, stop.
2. If the top entry of that column is zero, swap the top row with a row that has a nonzero entry in that column. This entry is the pivot.
3. For each row R_i below, apply $R_i \rightarrow R_i - \frac{a}{p} R_1$, where p is the pivot and a the entry of R_i in the pivot column. This makes every entry below the pivot zero.
4. Set the top row aside and repeat steps 1 to 3 on the rows below it, until no nonzero rows remain.

Scaling is never required, but dividing a row by its pivot (to make the pivot 1) often keeps the arithmetic simple.

Proof of Theorem 2.3.7. By induction on the number m of rows. A matrix with one row is in row echelon form. Let $m \geq 2$, and let A be an $m \times n$ matrix. If A is the zero matrix, it is in row echelon form. Otherwise let j be the leftmost column with a nonzero entry. After a swap (step 2) the entry $p = a_{1j}$ is nonzero, and after the replacements of step 3 every entry of column j below row 1 is zero. Columns $1, \dots, j-1$ were zero to begin with, and row operations keep them zero. So the matrix now has the form

$$\left[\begin{array}{c|c|c} 0 & p & * \\ 0 & 0 & B \end{array} \right],$$

where the first block column has $j-1$ columns and B has $m-1$ rows. By the induction hypothesis, row operations on the last $m-1$ rows bring B to row echelon form; these are also row operations on the whole matrix, and they leave row 1 and the zero columns unchanged. Now row 1 has its leading entry p in column j , every nonzero row below it has its leading entry in a column after j , and these leading entries step to the right with zero rows at the bottom because B is in row echelon form. So the whole matrix is in row echelon form, and it is row equivalent to A . $\square$

Example 2.3.8 (Forward elimination). Solve

$$\begin{aligned} x + 2y + z &= 3 \\ 2x + 5y + 3z &= 8 \\ -x + 2z &= 4. \end{aligned}$$

The augmented matrix has a 0 in row 3, column 2, because y does not appear in the third equation. The first pivot is the 1 in the top left corner. To clear the entries below it, we subtract $2R_1$ from R_2 ($2 - 2 \cdot 1 = 0$, $5 - 2 \cdot 2 = 1$, $3 - 2 \cdot 1 = 1$, $8 - 2 \cdot 3 = 2$) and add R_1 to R_3 ($-1 + 1 = 0$, $0 + 2 = 2$, $2 + 1 = 3$, $4 + 3 = 7$):

$$\left[\begin{array}{ccc|c} 1 & 2 & 1 & 3 \\ 2 & 5 & 3 & 8 \\ -1 & 0 & 2 & 4 \end{array} \right] \xrightarrow[\substack{R_2 \rightarrow R_2 - 2R_1 \\ R_3 \rightarrow R_3 + R_1}]{} \left[\begin{array}{ccc|c} 1 & 2 & 1 & 3 \\ 0 & 1 & 1 & 2 \\ 0 & 2 & 3 & 7 \end{array} \right].$$

Setting row 1 aside, the next pivot is the 1 in row 2, column 2. Subtracting $2R_2$ from R_3 ($0 - 0 = 0$, $2 - 2 \cdot 1 = 0$, $3 - 2 \cdot 1 = 1$, $7 - 2 \cdot 2 = 3$) gives

$$\xrightarrow{R_3 \rightarrow R_3 - 2R_2} \left[\begin{array}{ccc|c} 1 & 2 & 1 & 3 \\ 0 & 1 & 1 & 2 \\ 0 & 0 & 1 & 3 \end{array} \right],$$

which is in row echelon form, with pivots 1, 1, 1 on the diagonal. As equations it reads $x + 2y + z = 3$, $y + z = 2$, $z = 3$, and by Theorem 2.3.4 it has the same solutions as the original system.

2.3.5 Back substitution

A system in row echelon form is solved from the last nonzero equation upwards: each equation contains one new unknown, which is found from the values already known.

Example 2.3.9 (Back substitution). Continuing Example 2.3.8: the last equation gives $z = 3$. The second gives $y + 3 = 2$, so $y = -1$. The first gives $x + 2(-1) + 3 = 3$, that is $x + 1 = 3$, so $x = 2$. The system has exactly one solution,

$$(x, y, z) = (2, -1, 3).$$

Check in the original equations: $2 - 2 + 3 = 3$, $4 - 5 + 9 = 8$, $-2 + 6 = 4$. In the column picture of Section 2.2, the check says that

$$2 \begin{bmatrix} 1 \\ 2 \\ -1 \end{bmatrix} - \begin{bmatrix} 2 \\ 5 \\ 0 \end{bmatrix} + 3 \begin{bmatrix} 1 \\ 3 \\ 2 \end{bmatrix} = \begin{bmatrix} 3 \\ 8 \\ 4 \end{bmatrix} :$$

the right-hand side is a linear combination of the columns of the coefficient matrix, with the solution as weights.

Example 2.3.10 (A swap and a scaling). Solve $2y + 4z = 6$, $x + y + z = 2$, $2x + 3y + z = 1$. The top entry of the first column is 0, so it cannot be a pivot, and we swap the first two rows:

$$\left[\begin{array}{ccc|c} 0 & 2 & 4 & 6 \\ 1 & 1 & 1 & 2 \\ 2 & 3 & 1 & 1 \end{array} \right] \xrightarrow{R_1 \leftrightarrow R_2} \left[\begin{array}{ccc|c} 1 & 1 & 1 & 2 \\ 0 & 2 & 4 & 6 \\ 2 & 3 & 1 & 1 \end{array} \right] \xrightarrow{R_3 \rightarrow R_3 - 2R_1} \left[\begin{array}{ccc|c} 1 & 1 & 1 & 2 \\ 0 & 2 & 4 & 6 \\ 0 & 1 & -1 & -3 \end{array} \right].$$

($2 - 2 = 0$, $3 - 2 = 1$, $1 - 2 = -1$, $1 - 4 = -3$.) Scaling $R_2 \rightarrow \frac{1}{2}R_2$ makes the second pivot 1, and $R_3 \rightarrow R_3 - R_2$ clears the entry below it ($1 - 1 = 0$, $-1 - 2 = -3$, $-3 - 3 = -6$):

$$\xrightarrow{R_2 \rightarrow \frac{1}{2}R_2} \left[\begin{array}{ccc|c} 1 & 1 & 1 & 2 \\ 0 & 1 & 2 & 3 \\ 0 & 1 & -1 & -3 \end{array} \right] \xrightarrow{R_3 \rightarrow R_3 - R_2} \left[\begin{array}{ccc|c} 1 & 1 & 1 & 2 \\ 0 & 1 & 2 & 3 \\ 0 & 0 & -3 & -6 \end{array} \right].$$

Back substitution: $-3z = -6$ gives $z = 2$; $y + 2z = 3$ gives $y = 3 - 4 = -1$; $x + y + z = 2$ gives $x = 2 + 1 - 2 = 1$. The solution is $(x, y, z) = (1, -1, 2)$.

2.3.6 When there is no solution

Proposition 2.3.11. If a row echelon form of an augmented matrix $[A \mid \mathbf{b}]$ has a row $[0 \cdots 0 \mid c]$ with $c \neq 0$, then the system has no solution.

Proof. That row is the equation $0x_1 + \cdots + 0x_n = c$, that is $0 = c$, which no $\mathbf{x}$ satisfies. So the echelon system has no solution, and by Theorem 2.3.4 neither has the original system. $\square$

A system with no solution is called *inconsistent*. The converse of Proposition 2.3.11 is also true (if no such row appears, the system has a solution); it is proved in Section 2.5.

Example 2.3.12. For $x + y + 2z = 1$, $2x + 3y + 5z = 4$, $x + 2y + 3z = 6$:

$$\left[\begin{array}{ccc|c} 1 & 1 & 2 & 1 \\ 2 & 3 & 5 & 4 \\ 1 & 2 & 3 & 6 \end{array} \right] \xrightarrow[R_3 \rightarrow R_3 - R_1]{R_2 \rightarrow R_2 - 2R_1} \left[\begin{array}{ccc|c} 1 & 1 & 2 & 1 \\ 0 & 1 & 1 & 2 \\ 0 & 1 & 1 & 5 \end{array} \right] \xrightarrow{R_3 \rightarrow R_3 - R_2} \left[\begin{array}{ccc|c} 1 & 1 & 2 & 1 \\ 0 & 1 & 1 & 2 \\ 0 & 0 & 0 & 3 \end{array} \right].$$

The last row says $0 = 3$, so the system is inconsistent. If the third equation had been $x + 2y + 3z = 3$ instead, the same operations would end in the row $[0 \ 0 \ 0 \mid 3 - 1 - 2] = [0 \ 0 \ 0 \mid 0]$, the equation $0 = 0$. Then z could take any value, and the system would have infinitely many solutions. Such *free variables* are the subject of Sections 2.4 and 2.5.

2.3.7 Elimination over any field

Forward elimination and back substitution use only addition, subtraction, multiplication and division by nonzero numbers. Theorems 2.3.4 and 2.3.7 and their proofs use nothing else, so everything in this section holds word for word over any field $\mathbb{F}$ (Section 2.1): over $\mathbb{Q}$, over $\mathbb{C}$, and over the finite fields $\mathbb{Z}_p$.

Example 2.3.13 (Elimination in $\mathbb{Z}_5$). Solve $2x + y = 1$, $x + 4y = 2$ over $\mathbb{Z}_5$. The pivot is 2, and in $\mathbb{Z}_5$ we have $\frac{1}{2} = 3$ because $2 \cdot 3 = 6 = 1$. So the multiplier for row 2 is $1 \cdot 3 = 3$, and $R_2 \rightarrow R_2 - 3R_1$ gives $1 - 6 = -5 = 0$, $4 - 3 = 1$, $2 - 3 = -1 = 4$:

$$\left[\begin{array}{cc|c} 2 & 1 & 1 \\ 1 & 4 & 2 \end{array} \right] \xrightarrow{R_2 \rightarrow R_2 - 3R_1} \left[\begin{array}{cc|c} 2 & 1 & 1 \\ 0 & 1 & 4 \end{array} \right].$$

So $y = 4$, and $2x + 4 = 1$ gives $2x = -3 = 2$, so $x = 1$. Check: $2 + 4 = 6 = 1$ and $1 + 16 = 17 = 2$ in $\mathbb{Z}_5$.

Example 2.3.14 (The system of Section 2.2). Section 2.2 checked that $(1, 2, 3)$ solves system (2.1), $x + y + z = 6$, $2x - y + z = 3$, $x + 2y - z = 2$. Elimination shows it is the only solution. Clear the first column with $R_2 \rightarrow R_2 - 2R_1$ and $R_3 \rightarrow R_3 - R_1$; then swap the two lower rows, so that the next pivot is 1 and no fractions appear; then clear below it with $R_3 \rightarrow R_3 + 3R_2$:

$$\left[\begin{array}{ccc|c} 1 & 1 & 1 & 6 \\ 2 & -1 & 1 & 3 \\ 1 & 2 & -1 & 2 \end{array} \right] \rightarrow \left[\begin{array}{ccc|c} 1 & 1 & 1 & 6 \\ 0 & -3 & -1 & -9 \\ 0 & 1 & -2 & -4 \end{array} \right] \xrightarrow{R_2 \leftrightarrow R_3} \left[\begin{array}{ccc|c} 1 & 1 & 1 & 6 \\ 0 & 1 & -2 & -4 \\ 0 & -3 & -1 & -9 \end{array} \right] \\ \xrightarrow{R_3 \rightarrow R_3 + 3R_2} \left[\begin{array}{ccc|c} 1 & 1 & 1 & 6 \\ 0 & 1 & -2 & -4 \\ 0 & 0 & -7 & -21 \end{array} \right].$$

Every column of the coefficient part has a pivot, so back substitution leaves no choice: $-7z = -21$ gives $z = 3$; then $y - 6 = -4$ gives $y = 2$; then $x + 2 + 3 = 6$ gives $x = 1$. So $(1, 2, 3)$ is the unique solution, and the three planes of Section 2.2 meet in exactly one point. (Without the swap, $R_3 \rightarrow R_3 + \frac{1}{3}R_2$ gives the last row $[0 \ 0 \ -\frac{7}{3} \mid -7]$ and the same answer.)

Echelon forms are not unique: the matrix of Example 2.3.8 could equally have been reduced by first scaling or swapping rows, and the resulting staircase would have different entries. Section 2.4 goes one step further, to the *reduced* row echelon form, which is unique.

Check yourself

(a) Solve $x + y + z = 6$, $x + 2y + 3z = 14$, $2x + 3y + 5z = 23$ by Gaussian elimination, and check your answer.

(b) Is $\left[\begin{array}{ccc} 0 & 1 & 2 \\ 1 & 0 & 3 \end{array} \right]$ in row echelon form? If not, which single row operation brings it there?

Supplement: beyond the lecture

How much work is elimination? For a system of n equations in n unknowns, the first round of forward elimination updates $n - 1$ rows, each in n entries (the pivot column becomes zero without computation, and the right-hand side is included). The k -th round updates $n - k$ rows in $n - k + 1$ entries each. Each update is one multiplication and one subtraction, so forward elimination costs

$$\sum_{k=1}^{n-1} (n-k)(n-k+1) = \sum_{m=1}^{n-1} m(m+1) = \frac{(n-1)n(n+1)}{3} \approx \frac{n^3}{3}$$

multiply-subtract steps, while back substitution needs only about $n^2/2$. Doubling the size of a system multiplies the work by about eight. A laptop solves a system with a thousand unknowns in a fraction of a second; in Section 3.7 (Lecture 18) we will see that elimination is really a factorisation $A = LU$, which lets a computer reuse the work for many right-hand sides.

Pivoting and rounding. In exact arithmetic any nonzero entry can serve as a pivot. A computer, which rounds every result to about 16 significant digits, must choose more carefully. Consider $10^{-20}x + y = 1$, $x + y = 2$, whose solution is extremely close to $x = y = 1$. Using the tiny pivot 10^{-20} , the replacement $R_2 \rightarrow R_2 - 10^{20}R_1$ gives $(1 - 10^{20})y = 2 - 10^{20}$, and after rounding both sides equal -10^{20} , so $y = 1$; then $10^{-20}x = 1 - y = 0$ gives $x = 0$, which is completely wrong. Swapping the rows first, so that the pivot is 1, gives $(1 - 10^{-20})y = 1 - 2 \cdot 10^{-20}$, so $y = 1$ after rounding and then $x = 2 - y = 1$, which is correct to all 16 digits. Numerical software therefore uses *partial pivoting*: at each step it swaps up the entry of largest absolute value in the pivot column.

History. The eighth chapter of the *Nine Chapters on the Mathematical Art*, compiled in China over several centuries and complete by about the first century CE, is called *Fangcheng* (“rectangular arrays”). It sets out the coefficients of a system in columns of counting rods and eliminates unknowns by subtracting multiples of one column from another, exactly as we do with rows. Gauss used the method systematically in the early nineteenth century for his least-squares calculations of the orbit of the asteroid Pallas, and his name became attached to it much later.

2.4 Reduced row echelon form

Lecture 9

Gaussian elimination (Section 2.3) brings any matrix to a row echelon form, and back substitution then solves the system from the bottom up. Two things about echelon forms are unsatisfying. First, back substitution still leaves some algebra to do at the end. Second, echelon forms are not unique. For example, the echelon form we reached in Example 2.3.8 can be changed by one more replacement, $R_1 \rightarrow R_1 + R_2$ ($1 + 0 = 1$, $2 + 1 = 3$, $1 + 1 = 2$, $3 + 2 = 5$):

$$\left[\begin{array}{ccc|c} 1 & 2 & 1 & 3 \\ 0 & 1 & 1 & 2 \\ 0 & 0 & 1 & 3 \end{array} \right] \xrightarrow{R_1 \rightarrow R_1 + R_2} \left[\begin{array}{ccc|c} 1 & 3 & 2 & 5 \\ 0 & 1 & 1 & 2 \\ 0 & 0 & 1 & 3 \end{array} \right].$$

Both matrices are in row echelon form and, by Theorem 2.3.4, both describe the same solutions, but their entries differ, and neither is more correct than the other.

This section continues row reducing until there is nothing left to choose. The result, the *reduced* row echelon form, has two advantages: the solutions of the system can be read straight off it, and it is *unique*, so that the positions of the pivots, and the number of pivots, are properties of the matrix itself. Along the way we meet basic and free variables.

2.4.1 The definition

Definition 2.4.1 (Reduced row echelon form). A matrix is in *reduced row echelon form* if

- (i) all zero rows are at the bottom, and
 - (ii) each leading entry is strictly to the right of the leading entry of the row above it,
- that is, it is in row echelon form (Definition 2.3.5), and moreover
- (iii) every pivot equals 1, and
 - (iv) each pivot is the only nonzero entry in its column.

Conditions (i) and (ii) are those of a row echelon form; they already make every entry *below* a pivot zero. Condition (iv) adds that every entry *above* a pivot is zero too. Together with (iii) this says that each pivot column is as simple as a column can be: all zeros except for a single 1, and these 1s step down the staircase (Figure 2.7). The other columns may hold any numbers; conditions (iii) and (iv) say nothing about them.

1	*	0	0	*	*	*
0	0	1	0	*	*	*
0	0	0	1	*	*	*
0	0	0	0	0	0	0

1 = pivot
 0 = cleared above a pivot
 * = any number

Figure 2.7: The shape of a reduced row echelon form of a 4×7 augmented matrix (compare Figure 2.6). The shaded pivot columns contain a single 1 and zeros everywhere else.

Example 2.4.2. The matrix

$$\begin{bmatrix} 1 & 0 & 2 & 0 & 3 \\ 0 & 1 & -1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 4 \end{bmatrix}$$

is in reduced row echelon form: its pivots are the 1s in columns 1, 2 and 4, and every other entry of those three columns is zero. Columns 3 and 5 hold arbitrary numbers. Of the three matrices

$$\begin{bmatrix} 1 & 3 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 1 & 0 & 4 \\ 0 & 2 & 1 \end{bmatrix} \quad \begin{bmatrix} 1 & 5 & 0 \\ 0 & 1 & 2 \end{bmatrix}$$

the first is in reduced row echelon form. The second is not, because its second pivot is 2, not 1 (condition (iii)). The third is in row echelon form with pivots equal to 1, but it is not reduced: there is a 5 above its second pivot (condition (iv)).

2.4.2 The backward phase

To bring an echelon form into reduced form, we make each pivot 1 (by scaling its row) and clear the entries above it (by replacements with its row), working from the last pivot backwards.

Example 2.4.3 (Reducing an echelon form). Continue Example 2.3.8. Its echelon form already has pivots equal to 1. The last pivot is in row 3, column 3, with two 1s above it. The replacement $R_2 \rightarrow R_2 - R_3$

gives $0 - 0 = 0$, $1 - 0 = 1$, $1 - 1 = 0$, $2 - 3 = -1$, and $R_1 \rightarrow R_1 - R_3$ gives $1 - 0 = 1$, $2 - 0 = 2$, $1 - 1 = 0$, $3 - 3 = 0$:

$$\left[\begin{array}{ccc|c} 1 & 2 & 1 & 3 \\ 0 & 1 & 1 & 2 \\ 0 & 0 & 1 & 3 \end{array} \right] \xrightarrow[\substack{R_2 \rightarrow R_2 - R_3 \\ R_1 \rightarrow R_1 - R_3}]{\substack{R_2 \rightarrow R_2 - R_3 \\ R_1 \rightarrow R_1 - R_3}} \left[\begin{array}{ccc|c} 1 & 2 & 0 & 0 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & 3 \end{array} \right].$$

The next pivot is in row 2, column 2, with a 2 above it. The replacement $R_1 \rightarrow R_1 - 2R_2$ gives $1 - 2 \cdot 0 = 1$, $2 - 2 \cdot 1 = 0$, $0 - 2 \cdot 0 = 0$, $0 - 2 \cdot (-1) = 2$:

$$\xrightarrow{R_1 \rightarrow R_1 - 2R_2} \left[\begin{array}{ccc|c} 1 & 0 & 0 & 2 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & 3 \end{array} \right].$$

This is in reduced row echelon form. As equations it says $x = 2$, $y = -1$, $z = 3$: the solution $(2, -1, 3)$ of Example 2.3.9 stands in the last column, with no back substitution needed.

Working backwards is what keeps the cleared entries cleared. When the pivot row R_i is used to clear the entries above its pivot, the pivots below it have already been dealt with, so R_i has zeros in their columns (they were cleared earlier) and zeros to the left of its own pivot. Subtracting a multiple of R_i from a higher row therefore changes neither the zeros already made above later pivots nor anything to the left of the pivot of R_i .

2.4.3 Gauss–Jordan elimination

Gauss–Jordan elimination. Given a matrix:

1. *Forward phase.* Bring the matrix to row echelon form by forward elimination (Section 2.3).
2. *Backward phase.* For each pivot, starting with the last one and moving up:
 - (a) scale its row by $1/p$, where p is the pivot, so that the pivot becomes 1;
 - (b) for each row R_k above it, apply $R_k \rightarrow R_k - a R_i$, where R_i is the pivot row and a the entry of R_k in the pivot column. This makes every entry above the pivot zero.

Theorem 2.4.4 (Existence of the reduced row echelon form). Every matrix is row equivalent to a matrix in reduced row echelon form.

Proof. By Theorem 2.3.7, the matrix is row equivalent to a matrix E in row echelon form, with pivots in rows $1, \dots, r$ and columns $j_1 < j_2 < \dots < j_r$. We apply the backward phase to E and check, pivot by pivot, that the following statement stays true after pivot i has been treated (for $i = r, r-1, \dots, 1$):

the matrix is in row echelon form with the same pivot positions, and the pivots in rows $i, \dots, r$ equal 1 and are the only nonzero entries in their columns.

Before any step the statement holds with nothing to check about 1s (take $i = r+1$). Suppose it holds for $i+1$, and treat pivot i , in row i and column j_i .

Step (a) multiplies row i by $1/p \neq 0$. Entries that were zero stay zero, so the leading entry of row i stays in column j_i and becomes 1, and the zeros of row i in columns $j_{i+1}, \dots, j_r$ remain.

Step (b) replaces a row R_k with $k < i$ by $R_k - a R_i$. Row i is zero in columns $1, \dots, j_i - 1$ (its leading entry is in column j_i) and in columns $j_{i+1}, \dots, j_r$ (by the statement for $i+1$). So the replacement leaves

row k unchanged in all those columns: in particular its leading entry, in column $j_k < j_i$, does not move, and the zeros above the later pivots stay zero. In column j_i the new entry is $a - a \cdot 1 = 0$. Rows other than R_k are untouched.

So after pivot i the statement holds for i . After pivot 1 every pivot is 1 and alone in its column, and the matrix is still in row echelon form: it is in reduced row echelon form. All steps were row operations, so it is row equivalent to the original matrix. $\square$

Gauss–Jordan elimination uses only the four field operations, so, like Gaussian elimination, it works over any field $\mathbb{F}$. In Section 3.5 (Lecture 16) the same method computes the inverse of a matrix.

2.4.4 Basic and free variables

Example 2.4.5 (A free variable). At the end of Section 2.3 we met the system

$$x + y + 2z = 1, \quad 2x + 3y + 5z = 4, \quad x + 2y + 3z = 3.$$

The forward phase: $R_2 \rightarrow R_2 - 2R_1$ gives $[0 \ 1 \ 1 \mid 2]$; $R_3 \rightarrow R_3 - R_1$ gives $[0 \ 1 \ 1 \mid 2]$; then $R_3 \rightarrow R_3 - R_2$ gives a zero row:

$$\left[\begin{array}{ccc|c} 1 & 1 & 2 & 1 \\ 2 & 3 & 5 & 4 \\ 1 & 2 & 3 & 3 \end{array} \right] \xrightarrow[R_3 \rightarrow R_3 - R_1]{R_2 \rightarrow R_2 - 2R_1} \left[\begin{array}{ccc|c} 1 & 1 & 2 & 1 \\ 0 & 1 & 1 & 2 \\ 0 & 1 & 1 & 2 \end{array} \right] \xrightarrow{R_3 \rightarrow R_3 - R_2} \left[\begin{array}{ccc|c} 1 & 1 & 2 & 1 \\ 0 & 1 & 1 & 2 \\ 0 & 0 & 0 & 0 \end{array} \right].$$

The zero row is the equation $0 = 0$, which every (x, y, z) satisfies. In the backward phase the pivots, in columns 1 and 2, are already 1, and $R_1 \rightarrow R_1 - R_2$ clears the 1 above the second pivot ($1 - 0 = 1$, $1 - 1 = 0$, $2 - 1 = 1$, $1 - 2 = -1$):

$$\xrightarrow{R_1 \rightarrow R_1 - R_2} \left[\begin{array}{ccc|c} 1 & 0 & 1 & -1 \\ 0 & 1 & 1 & 2 \\ 0 & 0 & 0 & 0 \end{array} \right], \quad \text{that is} \quad \begin{array}{l} x + z = -1, \\ y + z = 2. \end{array}$$

The unknown z belongs to a column without a pivot, and nothing restricts it. Setting $z = t$, the two equations give $x = -1 - t$ and $y = 2 - t$, so the solutions are exactly

$$(x, y, z) = (-1 - t, 2 - t, t), \quad t \in \mathbb{R}.$$

Check with $t = 1$, that is $(-2, 1, 1)$: $-2 + 1 + 2 = 1$, $-4 + 3 + 5 = 4$, $-2 + 2 + 3 = 3$. Each value of t gives a different solution, so the system has infinitely many; geometrically, the three planes meet along a line.

Definition 2.4.6 (Basic and free variables). Consider a linear system whose augmented matrix $[A \mid \mathbf{b}]$ is row equivalent to the matrix $[R \mid \mathbf{d}]$ in reduced row echelon form. The unknown x_j is a *basic variable* if column j of R is a pivot column, and a *free variable* otherwise. Basic variables are also called *pivot variables*.

Only the columns before the bar belong to unknowns. A pivot in the last column (after the bar) means a row $[0 \ \cdots \ 0 \mid 1]$, the equation $0 = 1$, and then the system has no solution (Proposition 2.3.11). By Theorem 2.4.9 below, the reduced form is unique, so which variables are basic and which are free does not depend on how the system was reduced.

Proposition 2.4.7 (Reading off the solutions). Let $[R \mid \mathbf{d}]$ be in reduced row echelon form, with r nonzero rows and pivots in columns $j_1 < \cdots < j_r$, none of them the last column. Then

for every choice of values of the free variables there is exactly one solution of the system with those values, namely

$$x_{j_i} = d_i - \sum_{k \text{ free}} r_{ik} x_k \quad (i = 1, \dots, r).$$

Proof. Row i of $[R \mid \mathbf{d}]$ (for $i \leq r$) has a 1 in column j_i and zeros in all other pivot columns, by condition (iv). So its equation contains exactly one basic variable, x_{j_i} , with coefficient 1:

$$x_{j_i} + \sum_{k \text{ free}} r_{ik} x_k = d_i.$$

Rows $r+1, \dots, m$ are zero, including their last entry (otherwise the last column would hold a pivot), so they are the equation $0 = 0$. Hence, once the free variables are given values, each equation is satisfied exactly when its basic variable has the value displayed, and these values satisfy all the equations at once. $\square$

How to describe and use such solution sets in general (when a system has a solution, when it has exactly one, and how to write all solutions in vector form) is the subject of Section 2.5.

Example 2.4.8 (Two free variables). Solve

$$\begin{aligned} x_1 + 3x_2 + x_3 + x_4 &= 7 \\ 2x_1 + 6x_2 + 3x_3 + x_4 &= 16 \\ -x_1 - 3x_2 + x_3 - 3x_4 &= -3. \end{aligned}$$

Forward phase: $R_2 \rightarrow R_2 - 2R_1$ gives $[0 \ 0 \ 1 \ -1 \mid 2]$ ($2 - 2, 6 - 6, 3 - 2, 1 - 2, 16 - 14$), and $R_3 \rightarrow R_3 + R_1$ gives $[0 \ 0 \ 2 \ -2 \mid 4]$ ($-1 + 1, -3 + 3, 1 + 1, -3 + 1, -3 + 7$). Column 2 now has no nonzero entry below row 1, so the next pivot is in column 3, and $R_3 \rightarrow R_3 - 2R_2$ gives a zero row ($2 - 2 = 0, -2 + 2 = 0, 4 - 4 = 0$):

$$\left[\begin{array}{cccc|c} 1 & 3 & 1 & 1 & 7 \\ 2 & 6 & 3 & 1 & 16 \\ -1 & -3 & 1 & -3 & -3 \end{array} \right] \longrightarrow \left[\begin{array}{cccc|c} 1 & 3 & 1 & 1 & 7 \\ 0 & 0 & 1 & -1 & 2 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right] \xrightarrow{R_1 \rightarrow R_1 - R_2} \left[\begin{array}{cccc|c} 1 & 3 & 0 & 2 & 5 \\ 0 & 0 & 1 & -1 & 2 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right].$$

(The backward step: $1 - 0 = 1, 3 - 0 = 3, 1 - 1 = 0, 1 - (-1) = 2, 7 - 2 = 5$.) The pivot columns are 1 and 3, so x_1, x_3 are basic and x_2, x_4 are free. With $x_2 = s$ and $x_4 = t$, Proposition 2.4.7 gives $x_1 = 5 - 3s - 2t$ and $x_3 = 2 + t$:

$$(x_1, x_2, x_3, x_4) = (5 - 3s - 2t, s, 2 + t, t), \quad s, t \in \mathbb{R}.$$

Check with $s = 1, t = 0$, that is $(2, 1, 2, 0)$: $2 + 3 + 2 + 0 = 7, 4 + 6 + 6 + 0 = 16, -2 - 3 + 2 - 0 = -3$. Note that the pivot columns need not be consecutive: here column 2 is skipped.

2.4.5 Uniqueness of the reduced row echelon form

Theorem 2.4.9 (Uniqueness of the reduced row echelon form). Every matrix is row equivalent to *exactly one* matrix in reduced row echelon form.

Existence is Theorem 2.4.4. Before proving uniqueness, here it is at work.

Example 2.4.10 (Two paths, one reduced form). Reduce $A = \begin{bmatrix} 2 & 4 & 6 \\ 1 & 3 & 4 \end{bmatrix}$ in two ways.

Path 1. $R_1 \rightarrow \frac{1}{2}R_1$ gives $\begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 4 \end{bmatrix}$, and $R_2 \rightarrow R_2 - R_1$ gives $\begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \end{bmatrix}$; then $R_1 \rightarrow R_1 - 2R_2$ gives $\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$:

$$A \xrightarrow{R_1 \rightarrow \frac{1}{2}R_1} \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 4 \end{bmatrix} \xrightarrow{R_2 \rightarrow R_2 - R_1} \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \end{bmatrix} \xrightarrow{R_1 \rightarrow R_1 - 2R_2} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}.$$

Path 2. $R_1 \leftrightarrow R_2$, then $R_2 \rightarrow R_2 - 2R_1$ gives $\begin{bmatrix} 2 & 4 & 6 \\ 0 & -2 & -2 \end{bmatrix}$; then $R_2 \rightarrow -\frac{1}{2}R_2$ gives $\begin{bmatrix} 2 & 4 & 6 \\ 0 & 1 & 1 \end{bmatrix}$, and $R_1 \rightarrow R_1 - 2R_2$ gives $\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$:

$$A \xrightarrow{R_1 \leftrightarrow R_2} \begin{bmatrix} 1 & 3 & 4 \\ 2 & 4 & 6 \end{bmatrix} \xrightarrow{R_2 \rightarrow R_2 - 2R_1} \begin{bmatrix} 1 & 3 & 4 \\ 0 & -2 & -2 \end{bmatrix} \xrightarrow{R_2 \rightarrow -\frac{1}{2}R_2} \begin{bmatrix} 1 & 3 & 4 \\ 0 & 1 & 1 \end{bmatrix} \xrightarrow{R_1 \rightarrow R_1 - 3R_2} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}.$$

The two echelon forms $\begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \end{bmatrix}$ and $\begin{bmatrix} 1 & 3 & 4 \\ 0 & -2 & -2 \end{bmatrix}$ differ, but their pivots are in the same positions, and both paths end at the same reduced form.

The proof of uniqueness rests on one observation: row operations do not change the linear relations between the columns of a matrix. Write $\mathbf{a}_1, \dots, \mathbf{a}_n$ for the columns of an $m \times n$ matrix A (in this subsection, matrices are not augmented). By the column picture of Section 2.2, the homogeneous system with coefficient matrix A , that is the system whose right-hand sides are all 0, reads

$$x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \cdots + x_n\mathbf{a}_n = \mathbf{0}.$$

So a list of scalars $(c_1, \dots, c_n)$ is a solution of this system exactly when $c_1\mathbf{a}_1 + \cdots + c_n\mathbf{a}_n = \mathbf{0}$; we call such an equation a *column relation* of A .

Proposition 2.4.11 (Row operations keep column relations). If A and B are row equivalent $m \times n$ matrices with columns $\mathbf{a}_1, \dots, \mathbf{a}_n$ and $\mathbf{b}_1, \dots, \mathbf{b}_n$, then for all scalars $c_1, \dots, c_n$,

$$c_1\mathbf{a}_1 + \cdots + c_n\mathbf{a}_n = \mathbf{0} \iff c_1\mathbf{b}_1 + \cdots + c_n\mathbf{b}_n = \mathbf{0}.$$

Proof. Append a zero column to A and to B . The same row operations that turn A into B turn $[A \mid \mathbf{0}]$ into $[B \mid \mathbf{0}]$, since every row operation turns a zero entry of the last column into 0 again. By Theorem 2.3.4 the two homogeneous systems have the same solutions, and by the column picture these solutions are exactly the column relations. $\square$

Two consequences of row equivalence are used below. Since every row operation can be undone by a row operation (proof of Theorem 2.3.4), if A is row equivalent to B then B is row equivalent to A ; and if A is row equivalent to B and B to C , then A is row equivalent to C (perform one sequence of operations after the other). In particular, two matrices that are row equivalent to the same matrix are row equivalent to each other.

The second ingredient says that a matrix in reduced form is determined by its column relations. Write $\mathbf{e}_1, \dots, \mathbf{e}_m$ for the columns that have a single 1, in position 1, $\dots$, m respectively, and zeros elsewhere.

Proposition 2.4.12 (Columns of a reduced matrix). Let R be an $m \times n$ matrix in reduced row echelon form with pivots in columns $j_1 < \cdots < j_r$, and columns $\mathbf{r}_1, \dots, \mathbf{r}_n$. Then:

- (a) the pivot columns are $\mathbf{r}_{j_i} = \mathbf{e}_i$ for $i = 1, \dots, r$;
- (b) column k is a pivot column if and only if $\mathbf{r}_k$ is *not* a linear combination of $\mathbf{r}_1, \dots, \mathbf{r}_{k-1}$ (for

- $k = 1$: if and only if $\mathbf{r}_1 \neq \mathbf{0}$);
- (c) if column k is not a pivot column, then $\mathbf{r}_k = \sum_{i: j_i < k} r_{ik} \mathbf{r}_{j_i}$, where r_{ik} is the entry of R in row i and column k .

Proof. (a) The pivot of row i is a 1 in column j_i , and by condition (iv) every other entry of that column is 0.

(c) Let column k have no pivot, and let i be a row with $j_i > k$ or $i > r$. The leading entry of row i is to the right of column k (or row i is zero), so $r_{ik} = 0$. Hence the only possibly nonzero entries of $\mathbf{r}_k$ are r_{ik} for the rows i with $j_i < k$, and $\mathbf{r}_k = \sum_{i: j_i < k} r_{ik} \mathbf{e}_i = \sum_{i: j_i < k} r_{ik} \mathbf{r}_{j_i}$ by (a).

(b) If column k is not a pivot column, (c) writes $\mathbf{r}_k$ as a combination of earlier columns (the empty combination, $\mathbf{0}$, when no pivot lies to its left). Conversely let $k = j_i$ be a pivot column, so $\mathbf{r}_k = \mathbf{e}_i$. Every column $l < j_i$ has a zero in row i , since the leading entry of row i is in column j_i . So every combination of $\mathbf{r}_1, \dots, \mathbf{r}_{k-1}$ has a zero in row i and cannot equal $\mathbf{e}_i$. $\square$

Proof of Theorem 2.4.9. Let R and S be matrices in reduced row echelon form that are both row equivalent to A . Then R and S are row equivalent to each other, so by Proposition 2.4.11 they have exactly the same column relations.

Same pivot columns. Column k of R is a combination of the earlier columns, $\mathbf{r}_k = c_1 \mathbf{r}_1 + \dots + c_{k-1} \mathbf{r}_{k-1}$, exactly when $c_1 \mathbf{r}_1 + \dots + c_{k-1} \mathbf{r}_{k-1} - \mathbf{r}_k = \mathbf{0}$ is a column relation of R . That is also a column relation of S , so $\mathbf{s}_k$ is a combination of $\mathbf{s}_1, \dots, \mathbf{s}_{k-1}$ with the same coefficients, and conversely. By Proposition 2.4.12(b), R and S have the same pivot columns $j_1 < \dots < j_r$.

Same pivot column entries. By Proposition 2.4.12(a), $\mathbf{r}_{j_i} = \mathbf{e}_i = \mathbf{s}_{j_i}$ for every i .

Same other columns. Let column k have no pivot. By Proposition 2.4.12(c), $\mathbf{r}_k - \sum_{i: j_i < k} r_{ik} \mathbf{r}_{j_i} = \mathbf{0}$ is a column relation of R , hence of S : $\mathbf{s}_k = \sum_{i: j_i < k} r_{ik} \mathbf{s}_{j_i} = \sum_{i: j_i < k} r_{ik} \mathbf{e}_i = \mathbf{r}_k$.

So every column of S equals the corresponding column of R , and $R = S$. $\square$

In Example 2.4.10, the reduced form has third column $\mathbf{e}_1 + \mathbf{e}_2$, the sum of the two pivot columns, and indeed the columns of A satisfy the same relation: $(2, 1) + (4, 3) = (6, 4)$. The proof shows that this is no accident. Relations between columns are the central idea of Section 2.6 (Lecture 11), on linear independence.

Definition 2.4.13 (Pivot positions of a matrix). The *pivot positions* of a matrix A are the positions of the pivots (the leading 1s) in its reduced row echelon form, and the *pivot columns* of A are the columns that contain them.

Proposition 2.4.14. Every row echelon form of A has its pivots exactly in the pivot positions of A . In particular, all echelon forms of A have the same number of pivots.

Proof. Let E be a row echelon form of A . The backward phase of Gauss–Jordan elimination, applied to E , does not move any pivot (proof of Theorem 2.4.4) and ends in a reduced row echelon form row equivalent to A . By Theorem 2.4.9 that is *the* reduced form of A , so the pivots of E are in the pivot positions of A . $\square$

The number of pivots of a matrix is therefore a well-defined number attached to the matrix, independent of how it is reduced. In Section 5.2 it is called the *rank* of the matrix.

Check yourself

- (a) Find the reduced row echelon form of the augmented matrix of $x + 2y - z = 3$, $2x + 4y + z = 9$. Which variables are basic and which are free? Write down all solutions and check one.
- (b) True or false: if two matrices of the same size have the same reduced row echelon form, then they are row equivalent. Justify your answer.

Supplement: beyond the lecture

Is the backward phase worth it? For a single system with n equations, n unknowns and one solution, the backward phase of Gauss–Jordan elimination clears $n - 1$ entries above the last pivot, $n - 2$ above the one before, and so on. Each clearing updates only the entries to the right of the pivot column that are not already zero, which for a system in echelon form with all pivots 1 is just the right-hand side. So the backward phase costs about $n^2/2$ multiply-subtract steps, the same as back substitution, and both are small next to the $n^3/3$ of the forward phase (Section 2.3). If instead the backward phase is run on the full coefficient matrix before the forward phase has finished (clearing above and below each pivot as soon as it is found), the cost grows to about $n^3/2$, which is why computer programs keep the two phases separate.

Uniqueness depends on the field only through the arithmetic. The proof of Theorem 2.4.9 uses only the field axioms, so the reduced row echelon form is unique over any field. The same matrix of integers can, however, have different reduced forms over different fields, because the arithmetic differs. Over $\mathbb{Q}$, the matrix $\begin{bmatrix} 1 & 1 \\ 1 & 3 \end{bmatrix}$ reduces to $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ ($R_2 \rightarrow R_2 - R_1$ gives $[0 \ 2]$, then $R_2 \rightarrow \frac{1}{2}R_2$ and $R_1 \rightarrow R_1 - R_2$). Over $\mathbb{Z}_2$ the replacement gives $[0 \ 2] = [0 \ 0]$, so the reduced form is $\begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$, with only one pivot.

History. The backward phase is named after the German geodesist Wilhelm Jordan, who used it in 1888 in his handbook of surveying to solve the least-squares systems that arise in mapping land. The same method was described independently in the same year by B.-I. Clasen. The name *Gauss–Jordan* is therefore a little misleading twice over: the forward phase is far older than Gauss (Section 2.3), and the Jordan of the backward phase is not Camille Jordan, after whom the Jordan normal form is named.

2.5 Solution sets**Lecture 10**

Section 2.2 showed that a system of linear equations over $\mathbb{R}$ has no solution, exactly one, or infinitely many. Elimination (Sections 2.3 and 2.4) now lets us answer the three natural questions about any given system: *does a solution exist*, *is it unique*, and, when there are infinitely many, *what does the set of all solutions look like?* The answers can be read off an echelon form of the augmented matrix. Along the way we learn to describe infinitely many solutions with finitely many symbols, in parametric vector form, and we meet the homogeneous systems, whose right-hand sides are all zero. The section ends with the structure theorem: every solution set is one particular solution plus the solution set of the homogeneous system.

Throughout, $\mathbb{F}$ is a field, and the system has m equations in n unknowns $x_1, \dots, x_n$, with augmented matrix $[A \mid \mathbf{b}]$. Recall from Section 2.4 that the *pivot columns* of a matrix are the columns that hold the

pivots of any of its echelon forms (they are the same in all of them, Proposition 2.4.14), that an unknown whose column is a pivot column is a *basic variable*, also called a *pivot variable* (Definition 2.4.6), and that the other unknowns are the *free variables*.

2.5.1 Reading the solutions from the reduced form

Example 2.5.1. Solve the system of three equations in four unknowns

$$\begin{aligned}x_1 + 2x_2 - x_3 - x_4 &= 2 \\2x_1 + 4x_2 - x_3 &= 5 \\-x_1 - 2x_2 + 2x_3 + 3x_4 &= -1.\end{aligned}$$

The replacements $R_2 \rightarrow R_2 - 2R_1$ (that is, $2 - 2 = 0$, $4 - 4 = 0$, $-1 - (-2) = 1$, $0 - (-2) = 2$, $5 - 4 = 1$) and $R_3 \rightarrow R_3 + R_1$ ($-1 + 1 = 0$, $-2 + 2 = 0$, $2 - 1 = 1$, $3 - 1 = 2$, $-1 + 2 = 1$) clear the first column:

$$\left[\begin{array}{cccc|c} 1 & 2 & -1 & -1 & 2 \\ 2 & 4 & -1 & 0 & 5 \\ -1 & -2 & 2 & 3 & -1 \end{array} \right] \xrightarrow[R_3 \rightarrow R_3 + R_1]{R_2 \rightarrow R_2 - 2R_1} \left[\begin{array}{cccc|c} 1 & 2 & -1 & -1 & 2 \\ 0 & 0 & 1 & 2 & 1 \\ 0 & 0 & 1 & 2 & 1 \end{array} \right].$$

The second column has no pivot: below row 1 it holds only zeros. The next pivot is the 1 in row 2, column 3. Now $R_3 \rightarrow R_3 - R_2$ turns row 3 into a zero row, and $R_1 \rightarrow R_1 + R_2$ clears the entry above the pivot ($1 + 0 = 1$, $2 + 0 = 2$, $-1 + 1 = 0$, $-1 + 2 = 1$, $2 + 1 = 3$):

$$\xrightarrow[R_1 \rightarrow R_1 + R_2]{R_3 \rightarrow R_3 - R_2} \left[\begin{array}{cccc|c} 1 & 2 & 0 & 1 & 3 \\ 0 & 0 & 1 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right].$$

This is the reduced row echelon form. The pivot columns are columns 1 and 3, so x_1 and x_3 are pivot variables and x_2, x_4 are free. The last column is not a pivot column, and the zero row says only $0 = 0$. The two nonzero rows say

$$x_1 + 2x_2 + x_4 = 3, \quad x_3 + 2x_4 = 1,$$

and each of them contains exactly one pivot variable. Solving for the pivot variables,

$$x_1 = 3 - 2x_2 - x_4, \quad x_3 = 1 - 2x_4.$$

The free variables can take any values. Writing $x_2 = s$ and $x_4 = t$, the solutions are exactly the vectors

$$\mathbf{x} = \begin{bmatrix} 3 - 2s - t \\ s \\ 1 - 2t \\ t \end{bmatrix}, \quad s, t \in \mathbb{R}.$$

To separate the roles of s and t , split each entry into its constant part, its s part and its t part:

$$\mathbf{x} = \begin{bmatrix} 3 \\ 0 \\ 1 \\ 0 \end{bmatrix} + s \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} + t \begin{bmatrix} -1 \\ 0 \\ -2 \\ 1 \end{bmatrix} = \mathbf{p} + s\mathbf{v} + t\mathbf{w}.$$

For instance, $s = t = 1$ gives $\mathbf{x} = (3 - 2 - 1, 1, 1 - 2, 1) = (0, 1, -1, 1)$. Check in the original equations: $0 + 2 + 1 - 1 = 2$, $0 + 4 + 1 = 5$, $0 - 2 - 2 + 3 = -1$.

The formula of Example 2.5.1 describes every solution, and each solution exactly once. Indeed, by Theorem 2.3.4 a vector $\mathbf{x}$ solves the system exactly when it solves the reduced system, that is, when $x_1 = 3 - 2x_2 - x_4$ and $x_3 = 1 - 2x_4$. Such an $\mathbf{x}$ is $\mathbf{p} + s\mathbf{v} + t\mathbf{w}$ with $s = x_2$ and $t = x_4$, and these are the only values of s and t that produce it, because s and t are themselves entries of $\mathbf{x}$ (entries 2 and 4).

Definition 2.5.2 (General solution, parametric vector form). Let a consistent system have free variables $x_{j_1}, \dots, x_{j_k}$. Its *general solution* expresses every solution in terms of *parameters* $t_1, \dots, t_k \in \mathbb{F}$, the values given to the free variables. Written as

$$\mathbf{x} = \mathbf{p} + t_1 \mathbf{v}_1 + \dots + t_k \mathbf{v}_k,$$

with fixed vectors $\mathbf{p}, \mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{F}^n$, it is in *parametric vector form*. The vector $\mathbf{p}$ (all parameters zero) is itself a solution, called a *particular solution*.

So the solution set of Example 2.5.1 is

$$\{ \mathbf{p} + s\mathbf{v} + t\mathbf{w} : s, t \in \mathbb{R} \} = \mathbf{p} + \text{Span}\{\mathbf{v}, \mathbf{w}\},$$

the plane in $\mathbb{R}^4$ through the point $\mathbf{p}$ spanned by the directions $\mathbf{v}$ and $\mathbf{w}$ (here $\mathbf{p} + S$ means $\{\mathbf{p} + \mathbf{u} : \mathbf{u} \in S\}$). Two parameters, two free variables, a two-dimensional set of solutions.

The procedure always works in the same way. Bring the augmented matrix to reduced row echelon form; if it is consistent, move the free variables to the right-hand side of each nonzero row, rename them as parameters, and split the resulting vector into a constant part plus one vector per parameter. The vector multiplying the parameter of the free variable x_j has entry 1 in position j , entry 0 at every other free variable, and the negatives of the column- j entries of the reduced matrix at the pivot variables.

2.5.2 Existence and uniqueness

Theorem 2.5.3 (Existence and uniqueness). Let E be a row echelon form of the augmented matrix $[A \mid \mathbf{b}]$ of a system of m linear equations in n unknowns over a field $\mathbb{F}$.

- (a) (*Existence*) The system is consistent if and only if the last column of E is not a pivot column, that is, if and only if E has no row $[0 \cdots 0 \mid c]$ with $c \neq 0$.
- (b) (*Uniqueness*) Suppose the system is consistent, and let k be the number of free variables. Then for every choice of values of the free variables there is exactly one solution taking those values. In particular the solution is unique if and only if $k = 0$, that is, every column of A is a pivot column. If $k \geq 1$ and $\mathbb{F}$ is infinite (such as $\mathbb{R}$), there are infinitely many solutions.

Proof. By Theorem 2.3.4 the system and the echelon system of E have the same solutions, so we work with E . The pivot columns of E are those of $[A \mid \mathbf{b}]$.

(a) If E has a row $[0 \cdots 0 \mid c]$ with $c \neq 0$, the system is inconsistent by Proposition 2.3.11. In a row of that form the leading entry is c , in the last column, so such a row exists exactly when the last column of E is a pivot column. Conversely, suppose the last column is not a pivot column. Then every nonzero row of E has its pivot among the first n columns, so it is an equation whose first nonzero coefficient multiplies a pivot variable; every zero row is the equation $0 = 0$, which always holds. Give the free variables any values. We now find the pivot variables from the bottom nonzero row upwards. If the pivot of row r is $p_r \neq 0$ in column j_r , the row reads

$$p_r x_{j_r} + (\text{terms in the unknowns } x_j \text{ with } j > j_r) = d_r,$$

because all entries of row r before column j_r are zero. Every unknown x_j with $j > j_r$ is either free (already chosen) or the pivot variable of a lower row (already found, because pivots step to the right). So the row is one equation $p_r x_{j_r} = d_r - (\text{known terms})$ in the single unknown x_{j_r} , and since $p_r \neq 0$

it has exactly one solution $x_{j_r} = p_r^{-1}(d_r - (\text{known terms}))$ by Proposition 2.1.6. Working up through all the nonzero rows we obtain values for all unknowns that satisfy every row of E . So the system is consistent.

(b) The construction in (a) starts from arbitrary values of the free variables and, at each step, *must* take the value it takes, since each step solves an equation with exactly one solution. So for each choice of values of the free variables there is exactly one solution with those values; and every solution arises in this way, from the values of its own free entries. Different choices give different solutions, because they differ in some entry. Hence the solutions correspond one-to-one to the lists of k values in $\mathbb{F}$. If $k = 0$ there is just one (empty) list, and the solution is unique. If $k \geq 1$, a single free variable already takes every value in $\mathbb{F}$, so when $\mathbb{F}$ is infinite there are infinitely many solutions. $\square$

Over $\mathbb{R}$ this refines Theorem 2.2.7: the echelon form tells us *which* of the three cases occurs. The test is a test on columns.

last column is a pivot column	no solution
last column is not; every column of A is a pivot column	exactly one solution
last column is not; some column of A is not	infinitely many solutions (over $\mathbb{R}$)

Example 2.5.4. Three systems in the unknowns x_1, x_2, x_3 over $\mathbb{R}$ have the echelon forms

$$\left[\begin{array}{ccc|c} 1 & 2 & -1 & 4 \\ 0 & 1 & 3 & 2 \\ 0 & 0 & 0 & 5 \end{array} \right], \quad \left[\begin{array}{ccc|c} 2 & 1 & 0 & 1 \\ 0 & 3 & 1 & 0 \\ 0 & 0 & 4 & 8 \end{array} \right], \quad \left[\begin{array}{ccc|c} 1 & -1 & 2 & 3 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{array} \right].$$

The first has a pivot in the last column (the row $[0 \ 0 \ 0 \mid 5]$), so it has no solution. In the second the last column is not a pivot column and all three columns of A are, so it has exactly one solution. In the third the last column is not a pivot column, but column 2 of A is not either: x_2 is free, and there are infinitely many solutions. No back substitution was needed to decide this.

Each row of an echelon form holds at most one pivot, so a matrix with m rows has at most m pivot columns. This simple count has an important consequence.

Proposition 2.5.5 (More unknowns than equations). If a consistent system of m equations in n unknowns has $n > m$, then it has at least $n - m$ free variables. Over $\mathbb{R}$ (or any infinite field) it therefore has infinitely many solutions.

Proof. At most m of the n columns of A are pivot columns, so at least $n - m \geq 1$ of them are not. Apply Theorem 2.5.3(b). $\square$

Note the word *consistent*: a system with more unknowns than equations can still have no solution, as $x_1 + x_2 + x_3 = 0$, $x_1 + x_2 + x_3 = 1$ shows. And a system with *more* equations than unknowns can have exactly one solution, or infinitely many: counting equations alone never settles the question; the pivots do.

Remark 2.5.6 (Finite fields). Theorem 2.5.3(b) also counts solutions over a finite field: a consistent system over $\mathbb{Z}_p$ (p prime) with k free variables has exactly p^k solutions, one for each list of k values in $\mathbb{Z}_p$. For example, over $\mathbb{Z}_3$ the single equation $x_1 + x_2 = 2$ has one free variable, x_2 , and $3^1 = 3$ solutions: $(2, 0)$, $(1, 1)$ and $(0, 2)$, since $2 + 0 = 2$, $1 + 1 = 2$ and $0 + 2 = 2$ in $\mathbb{Z}_3$. This is why “no solution, one, or infinitely many” needs an infinite field (Remark 2.2.8).

2.5.3 Homogeneous systems

Definition 2.5.7 (Homogeneous system). A system of linear equations is *homogeneous* if all its right-hand sides are zero, that is, its augmented matrix is $[A \mid \mathbf{0}]$. The zero vector $\mathbf{x} = \mathbf{0}$ is always a solution, called the *trivial solution*; any other solution is *nontrivial*.

The trivial solution works because each equation reads $a_{i1} \cdot 0 + \cdots + a_{in} \cdot 0 = 0$, and $a \cdot 0 = 0$ for every $a \in \mathbb{F}$ (Proposition 2.1.5(a), with (M2)). So a homogeneous system is *always* consistent, and the only question is whether it has solutions other than $\mathbf{0}$.

Proposition 2.5.8. Row operations keep the last column of $[A \mid \mathbf{0}]$ zero. Consequently a homogeneous system has a nontrivial solution if and only if it has at least one free variable.

Proof. A swap only moves the zeros of the last column. A scaling replaces an entry 0 by $c \cdot 0 = 0$ (Proposition 2.1.5(a), with (M2)). A replacement $R_i \rightarrow R_i + c R_j$ replaces it by $0 + c \cdot 0 = 0 + 0 = 0$ (the same rule, then (A3)). So every echelon form of $[A \mid \mathbf{0}]$ has a zero last column; in particular the last column is not a pivot column, as Theorem 2.5.3(a) requires.

If there are no free variables, the solution is unique by Theorem 2.5.3(b), and since $\mathbf{0}$ is a solution it is the only one. If x_j is free, Theorem 2.5.3(b) gives a solution with $x_j = 1$, and that solution is not $\mathbf{0}$. $\square$

Because the last column stays zero, it is usual to row reduce only the coefficient matrix A of a homogeneous system.

Theorem 2.5.9 (Nontrivial solutions). A homogeneous system with more unknowns than equations ($n > m$) has a nontrivial solution. This holds over every field.

Proof. The system is consistent, because $\mathbf{0}$ is a solution, so by Proposition 2.5.5 it has at least $n - m \geq 1$ free variables, and by Proposition 2.5.8 it has a nontrivial solution. $\square$

Theorem 2.5.9 looks modest, but it is one of the workhorses of the course. In Section 2.6 it shows that too many vectors are always linearly dependent, and in Chapter 4 it is the key step in proving that every basis of a space has the same number of vectors.

Example 2.5.10. The homogeneous system with the coefficients of Example 2.5.1,

$$x_1 + 2x_2 - x_3 - x_4 = 0, \quad 2x_1 + 4x_2 - x_3 = 0, \quad -x_1 - 2x_2 + 2x_3 + 3x_4 = 0,$$

is reduced by the same row operations, with zeros after the bar throughout, to

$$\left[\begin{array}{cccc|c} 1 & 2 & 0 & 1 & 0 \\ 0 & 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right].$$

So $x_1 = -2x_2 - x_4$ and $x_3 = -2x_4$, and with $x_2 = s$, $x_4 = t$,

$$\mathbf{x} = s \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} + t \begin{bmatrix} -1 \\ 0 \\ -2 \\ 1 \end{bmatrix} = s \mathbf{v} + t \mathbf{w}.$$

The solution set is $\text{Span}\{\mathbf{v}, \mathbf{w}\}$, a plane through the origin of $\mathbb{R}^4$. The same vectors $\mathbf{v}, \mathbf{w}$ appeared in Example 2.5.1; only the particular solution $\mathbf{p}$ has gone. The nontrivial solutions are guaranteed in advance by Theorem 2.5.9, since $n = 4 > 3 = m$.

In general the solution set of a homogeneous system is the span of the vectors in its parametric vector form, so it always contains $\mathbf{0}$ and is closed under adding solutions and scaling them.

2.5.4 Particular plus homogeneous

Comparing Examples 2.5.1 and 2.5.10 suggests the general rule. As in Section 2.2, write $\mathbf{r}_i = (a_{i1}, \dots, a_{in})$ for the row of coefficients of equation i , so that the equation reads $\mathbf{r}_i \cdot \mathbf{x} = b_i$. The rules of the dot product hold over any field (Section 2.2).

Theorem 2.5.11 (Structure of the solution set). Let $\mathbf{p}$ be one solution of a system with augmented matrix $[A \mid \mathbf{b}]$, and let H be the solution set of the homogeneous system $[A \mid \mathbf{0}]$. Then the solution set of $[A \mid \mathbf{b}]$ is

$$\mathbf{p} + H = \{ \mathbf{p} + \mathbf{h} : \mathbf{h} \in H \}.$$

In words: every solution is the particular solution $\mathbf{p}$ plus a solution of the homogeneous system, and every such sum is a solution.

Proof. We show two inclusions. Fix an equation i ; we know $\mathbf{r}_i \cdot \mathbf{p} = b_i$.

Every $\mathbf{p} + \mathbf{h}$ is a solution. Let $\mathbf{h} \in H$, so $\mathbf{r}_i \cdot \mathbf{h} = 0$. By additivity of the dot product (Theorem 1.4.6(b)),

$$\mathbf{r}_i \cdot (\mathbf{p} + \mathbf{h}) = \mathbf{r}_i \cdot \mathbf{p} + \mathbf{r}_i \cdot \mathbf{h} = b_i + 0 = b_i,$$

the last step by (A3). This holds for every i , so $\mathbf{p} + \mathbf{h}$ is a solution.

Every solution has this form. Let $\mathbf{x}$ be a solution, so $\mathbf{r}_i \cdot \mathbf{x} = b_i$, and put $\mathbf{h} = \mathbf{x} - \mathbf{p} = \mathbf{x} + (-1)\mathbf{p}$. Then $\mathbf{x} = \mathbf{p} + \mathbf{h}$ by the rules of vector algebra (Theorem 1.2.7, which hold in $\mathbb{F}^n$). By additivity (Theorem 1.4.6(b)), then symmetry and homogeneity (Theorem 1.4.6(a) and (c)) to move the scalar -1 out of the second factor,

$$\mathbf{r}_i \cdot \mathbf{h} = \mathbf{r}_i \cdot \mathbf{x} + \mathbf{r}_i \cdot ((-1)\mathbf{p}) = \mathbf{r}_i \cdot \mathbf{x} + (-1)(\mathbf{r}_i \cdot \mathbf{p}) = b_i + (-1)b_i = b_i - b_i = 0,$$

using Proposition 2.1.5(b) for $(-1)b_i = -b_i$ and (A4) for the last step. This holds for every i , so $\mathbf{h} \in H$ and $\mathbf{x} \in \mathbf{p} + H$. $\square$

Notice that the theorem works with *any* particular solution. In Example 2.5.1, $\mathbf{p} = (3, 0, 1, 0)$ came from setting both parameters to zero, but the solution $(0, 1, -1, 1)$ would serve equally well: it differs from $\mathbf{p}$ by $\mathbf{v} + \mathbf{w} \in H$, so it describes the same set.

Example 2.5.12 (A line and its parallel through the origin). In $\mathbb{R}^2$, the equation $x + 2y = 4$ is a system of one equation in two unknowns. Its homogeneous version $x + 2y = 0$ has the free variable $y = t$ and the solutions $t(-2, 1)$: the line H through the origin in the direction $(-2, 1)$. One particular solution of $x + 2y = 4$ is $\mathbf{p} = (4, 0)$. By Theorem 2.5.11 the solution set is

$$\mathbf{p} + H = \{ (4, 0) + t(-2, 1) : t \in \mathbb{R} \},$$

the line through $(4, 0)$ parallel to H (Figure 2.8). Another particular solution, $\mathbf{q} = (2, 1)$, gives the same line, since $\mathbf{q} = \mathbf{p} + 1 \cdot (-2, 1)$.




Figure 2.8: The solution set of $x + 2y = 4$ (green) is the solution set H of $x + 2y = 0$ (blue), moved by the particular solution $\mathbf{p}$. Any other particular solution, such as $\mathbf{q}$, gives the same line.

So the solution set of a consistent system always has the same shape as that of its homogeneous system: a point (when $H = \{\mathbf{0}\}$), a line, a plane, or a higher-dimensional flat piece, moved away from the origin by a particular solution. Unless $\mathbf{b} = \mathbf{0}$, it does not pass through the origin, and so it is not a span. The number of parameters, which is the number of free variables, measures how large the solution set is; Chapter 4 makes this precise with the notion of dimension.

Check yourself

- (a) A system in the unknowns x_1, x_2, x_3 has the reduced row echelon form $\left[\begin{array}{ccc|c} 1 & 0 & 3 & 2 \\ 0 & 1 & -1 & 5 \\ 0 & 0 & 0 & 0 \end{array} \right]$.

Does it have no solution, one, or infinitely many? Write its general solution in parametric vector form, and its homogeneous system's solution set as a span.

- (b) Without row reducing, explain why a homogeneous system of 3 equations in 5 unknowns has a nontrivial solution. Must every system of 3 equations in 5 unknowns have a solution?

Supplement: beyond the lecture

Different parameters, same set. A parametric vector form is not unique. In Example 2.5.1 we could have chosen $x_2 = s$, $x_4 = t$ as above, or used any other particular solution (Theorem 2.5.11), or combinations such as $\mathbf{v} + \mathbf{w}$ and $\mathbf{w}$ as the two direction vectors, since $\text{Span}\{\mathbf{v} + \mathbf{w}, \mathbf{w}\} = \text{Span}\{\mathbf{v}, \mathbf{w}\}$. What does not change is the number of parameters: Section 2.4 showed that the pivot columns, and hence the number of free variables, do not depend on how we row reduce. In Chapter 6 this number becomes the dimension of the kernel, and the rank-nullity theorem says that the number of pivot variables plus the number of free variables is n .

Balancing a chemical equation. When methane burns, a molecules of CH_4 and b molecules of O_2 become c molecules of CO_2 and d molecules of H_2O . Counting the atoms of each element on both sides gives a homogeneous system:

$$\text{carbon: } a - c = 0, \quad \text{hydrogen: } 4a - 2d = 0, \quad \text{oxygen: } 2b - 2c - d = 0.$$

It has 4 unknowns and 3 equations, so by Theorem 2.5.9 a nontrivial solution exists. The equations give $c = a$, $d = 2a$, and $2b = 2a + 2a$, so $b = 2a$: one free variable, and the solutions $(a, b, c, d) = a(1, 2, 1, 2)$.

The smallest whole-number choice $a = 1$ is the balanced equation $\text{CH}_4 + 2 \text{O}_2 \rightarrow \text{CO}_2 + 2 \text{H}_2\text{O}$. Check: one carbon, four hydrogen and four oxygen atoms on each side.

Superposition, once more. Theorem 2.5.11 is the superposition principle of Section 2.2 in another form. Engineers solving a linear model (an electrical network, a loaded structure) find one response to the given forces and add every response to zero force; the second part is often called the *free response* of the system.

2.6 Linear independence

Lecture 11

In Section 1.3 we met the span of a list of vectors: everything that can be reached by linear combinations of them. This section asks a sharper question: does every vector in the list really contribute something? Take

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \mathbf{v}_3 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}.$$

The first two span the horizontal plane of vectors with third entry 0, and $\mathbf{v}_3 = \mathbf{v}_1 + \mathbf{v}_2$ lies in that plane too ($1 + 0 = 1$, $0 + 1 = 1$, $0 + 0 = 0$). By Theorem 1.3.11, adding $\mathbf{v}_3$ does not enlarge the span: three vectors, but only two different directions. If instead the third vector is $\mathbf{u} = (0, 0, 1)$, which points straight up, then none of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{u}$ lies in the plane of the other two, and together they span all of $\mathbb{R}^3$ (Example 1.3.9). Lists without a redundant vector are called *linearly independent*. They are the building blocks of bases and dimension (Sections 4.2 and 4.5, Lectures 20 and 23).

Throughout, $\mathbb{F}$ is a field and the vectors lie in $\mathbb{F}^n$ (Definition 2.1.18); if no field is named, $\mathbb{F} = \mathbb{R}$. The proofs use only the rules of vector algebra (Theorem 1.2.7) and their consequences (Proposition 1.2.8), which hold in $\mathbb{F}^n$ for every field.

2.6.1 The definition

Which vector of a list is the redundant one? In the example, $\mathbf{v}_3 = \mathbf{v}_1 + \mathbf{v}_2$, but equally $\mathbf{v}_1 = \mathbf{v}_3 - \mathbf{v}_2$. To avoid choosing, we move everything to one side:

$$1 \mathbf{v}_1 + 1 \mathbf{v}_2 + (-1) \mathbf{v}_3 = \mathbf{0}.$$

This is a linear combination of all three vectors, in which no vector plays a special role. There is always one way to obtain $\mathbf{0}$: give every vector the weight 0, since $0 \mathbf{v}_1 + \cdots + 0 \mathbf{v}_k = \mathbf{0} + \cdots + \mathbf{0} = \mathbf{0}$ by Proposition 1.2.8(d) and rule (3). This is the *trivial combination*, and it says nothing about the vectors.

Definition 2.6.1 (Linear independence). Vectors $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{F}^n$ are *linearly independent* if the only scalars $c_1, \dots, c_k \in \mathbb{F}$ with

$$c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2 + \cdots + c_k \mathbf{v}_k = \mathbf{0}$$

are $c_1 = c_2 = \cdots = c_k = 0$. They are *linearly dependent* if there are scalars $c_1, \dots, c_k$, *not all zero*, with $c_1 \mathbf{v}_1 + \cdots + c_k \mathbf{v}_k = \mathbf{0}$; such an equation is called a *linear dependence relation*.

Every list is either independent or dependent, and not both. We speak of a list $\mathbf{v}_1, \dots, \mathbf{v}_k$ rather than a set, because the same vector may occur twice; a list in which $\mathbf{v}_i = \mathbf{v}_j$ for some $i \neq j$ is dependent, since

$\mathbf{v}_i - \mathbf{v}_j = \mathbf{0}$ is a relation with weights 1 and -1 . The words “the set $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is independent” are used with the same meaning when the $\mathbf{v}_i$ are distinct.

Example 2.6.2. With $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ as above, $\mathbf{v}_1 + \mathbf{v}_2 - \mathbf{v}_3 = \mathbf{0}$ is a dependence relation, so $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ are dependent.

To show that $\mathbf{v}_1, \mathbf{v}_2, \mathbf{u}$ are independent, suppose $c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + c_3\mathbf{u} = \mathbf{0}$. Entry by entry, the left side is

$$\begin{bmatrix} c_1 \cdot 1 + c_2 \cdot 0 + c_3 \cdot 0 \\ c_1 \cdot 0 + c_2 \cdot 1 + c_3 \cdot 0 \\ c_1 \cdot 0 + c_2 \cdot 0 + c_3 \cdot 1 \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix},$$

so the equation says $c_1 = 0$, $c_2 = 0$ and $c_3 = 0$. Only the trivial combination gives $\mathbf{0}$, so the three vectors are independent. The same argument shows that the standard vectors $\mathbf{e}_1, \dots, \mathbf{e}_n$ of $\mathbb{F}^n$ are independent.

2.6.2 Dependent means that one vector is redundant

Theorem 2.6.3. Let $k \geq 2$. Vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ are linearly dependent if and only if at least one of them is a linear combination of the others.

Proof. By rules (1) and (2), the terms of a linear combination may be reordered freely, so we may number the vectors as we like.

($\Leftarrow$) Suppose one of the vectors is a combination of the others, and number the vectors so that it is the last one:

$$\mathbf{v}_k = d_1\mathbf{v}_1 + \dots + d_{k-1}\mathbf{v}_{k-1}.$$

Add $-\mathbf{v}_k$ to both sides. On the left, $\mathbf{v}_k + (-\mathbf{v}_k) = \mathbf{0}$ by rule (4). On the right, $-\mathbf{v}_k = (-1)\mathbf{v}_k$ by Proposition 1.2.8(f). Hence

$$d_1\mathbf{v}_1 + \dots + d_{k-1}\mathbf{v}_{k-1} + (-1)\mathbf{v}_k = \mathbf{0}.$$

The weight of $\mathbf{v}_k$ is $-1 \neq 0$, so this is a dependence relation, and the vectors are dependent.

($\Rightarrow$) Suppose $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k = \mathbf{0}$ with not all c_i zero, and number the vectors so that $c_k \neq 0$. Add $(-c_1)\mathbf{v}_1 + \dots + (-c_{k-1})\mathbf{v}_{k-1}$ to both sides. For each $i < k$, rule (8) and Proposition 1.2.8(d) give $c_i\mathbf{v}_i + (-c_i)\mathbf{v}_i = (c_i - c_i)\mathbf{v}_i = 0\mathbf{v}_i = \mathbf{0}$, so with rules (1) to (3) the left side becomes $c_k\mathbf{v}_k$:

$$c_k\mathbf{v}_k = (-c_1)\mathbf{v}_1 + \dots + (-c_{k-1})\mathbf{v}_{k-1}.$$

Since $c_k \neq 0$, it has an inverse c_k^{-1} in $\mathbb{F}$. Multiply both sides by c_k^{-1} . On the left, rules (6) and (5) give $c_k^{-1}(c_k\mathbf{v}_k) = (c_k^{-1}c_k)\mathbf{v}_k = 1\mathbf{v}_k = \mathbf{v}_k$. On the right, rule (7) distributes c_k^{-1} over the sum, and rule (6) gives $c_k^{-1}((-c_i)\mathbf{v}_i) = (-c_i/c_k)\mathbf{v}_i$. Therefore

$$\mathbf{v}_k = \left(-\frac{c_1}{c_k}\right)\mathbf{v}_1 + \dots + \left(-\frac{c_{k-1}}{c_k}\right)\mathbf{v}_{k-1},$$

a linear combination of the others. □

The proof divides by the weight c_k , which is possible only because $c_k \neq 0$; this is why the definition asks for weights that are *not all zero*. Note also what the theorem does not say: in a dependent list *some* vector is a combination of the others, not necessarily every one.

Example 2.6.4. The vectors $(1, 0)$, $(2, 0)$, $(0, 1)$ are dependent: $2(1, 0) - (2, 0) + 0(0, 1) = (0, 0)$. The second vector is $2(1, 0) + 0(0, 1)$, and the first is $\frac{1}{2}(2, 0) + 0(0, 1)$. But $(0, 1)$ is not a combination of the other two, since every combination $a(1, 0) + b(2, 0) = (a + 2b, 0)$ has second entry 0. In the relation above, the weight of $(0, 1)$ is 0, so the proof of Theorem 2.6.3 cannot solve for it.

2.6.3 Small lists

Proposition 2.6.5. Let $\mathbf{v}, \mathbf{w}, \mathbf{v}_1, \dots, \mathbf{v}_k, \mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{F}^n$.

- (a) A single vector $\mathbf{v}$ is independent if and only if $\mathbf{v} \neq \mathbf{0}$.
- (b) Every list that contains $\mathbf{0}$ is dependent.
- (c) Two vectors $\mathbf{v}, \mathbf{w}$ are dependent if and only if one of them is a scalar multiple of the other.
- (d) If $\mathbf{v}_1, \dots, \mathbf{v}_k$ are dependent, then so are $\mathbf{v}_1, \dots, \mathbf{v}_k, \mathbf{u}_1, \dots, \mathbf{u}_m$. Equivalently, if a list is independent, so is every list obtained from it by deleting vectors.

Proof. (a) If $\mathbf{v} \neq \mathbf{0}$ and $c\mathbf{v} = \mathbf{0}$, then $c = 0$ by Proposition 1.2.8(g); so $\mathbf{v}$ is independent. If $\mathbf{v} = \mathbf{0}$, then $1 \cdot \mathbf{0} = \mathbf{0}$ (rule (5)) is a relation with the weight $1 \neq 0$, so $\mathbf{0}$ is dependent.

(b) Suppose $\mathbf{v}_j = \mathbf{0}$. Give $\mathbf{v}_j$ the weight 1 and every other vector the weight 0. By rule (5) and Proposition 1.2.8(d) every term is $\mathbf{0}$, so the combination is $\mathbf{0}$ by rule (3), and the weight 1 is not zero.

(c) This is Theorem 2.6.3 with $k = 2$: “a combination of the other one” means a scalar multiple of it.

(d) Take a relation $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k = \mathbf{0}$ with some $c_i \neq 0$, and give each $\mathbf{u}_j$ the weight 0. The new terms are $0\mathbf{u}_j = \mathbf{0}$, so the longer combination is still $\mathbf{0}$ (rule (3)), and c_i is still a nonzero weight. The second sentence is the contrapositive: if a shorter list obtained by deleting vectors were dependent, the longer list would be dependent too. $\square$

Geometrically, (c) says that two vectors are dependent exactly when they lie on one line through the origin (Theorem 1.3.8 for $\mathbb{R}^2$).

Example 2.6.6. $(1, 2)$ and $(2, 4)$ are dependent: $2(1, 2) - (2, 4) = (2 \cdot 1 - 2, 2 \cdot 2 - 4) = (0, 0)$. On the other hand, $(1, 2)$ and $(3, 1)$ are independent. Neither is a multiple of the other: $c(1, 2) = (3, 1)$ would need $c = 3$ from the first entry and then $2c = 6 \neq 1$; and $c(3, 1) = (1, 2)$ would need $c = \frac{1}{3}$ and then $c = 2$. By (c), they are independent.

2.6.4 The geometric meaning

In $\mathbb{R}^2$, two vectors are independent when they point along different lines through the origin; then by Theorem 1.3.8 their span is all of $\mathbb{R}^2$, so any third vector of $\mathbb{R}^2$ lies in their span and three vectors of $\mathbb{R}^2$ are always dependent (this also follows from Theorem 2.6.10 below). In $\mathbb{R}^3$, three vectors are dependent when they lie in one plane through the origin, like $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ above, and independent when they do not (Figure 2.9).

A good way to see independence is to add the vectors one at a time and watch the span grow: the first vector must not be $\mathbf{0}$, so it spans a line; the second must leave that line, so the two span a plane; the third must leave that plane; and so on.

Theorem 2.6.7. Vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ are linearly independent if and only if $\mathbf{v}_1 \neq \mathbf{0}$ and, for each $j = 2, \dots, k$,

$$\mathbf{v}_j \notin \text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_{j-1}\}.$$

Proof. ($\Rightarrow$) We prove the contrapositive. If $\mathbf{v}_1 = \mathbf{0}$, the list is dependent by Proposition 2.6.5(b). If $\mathbf{v}_j = d_1\mathbf{v}_1 + \dots + d_{j-1}\mathbf{v}_{j-1}$ for some $j \geq 2$, then, exactly as in the first half of the proof of Theorem 2.6.3,

$$d_1\mathbf{v}_1 + \dots + d_{j-1}\mathbf{v}_{j-1} + (-1)\mathbf{v}_j + 0\mathbf{v}_{j+1} + \dots + 0\mathbf{v}_k = \mathbf{0}$$




Figure 2.9: Left: $\mathbf{v} = (2, 1)$ and $\mathbf{w} = (1, 2)$ point along different lines and are independent; $\mathbf{u} = \mathbf{v} + \mathbf{w}$ lies in their span, so $\mathbf{v}, \mathbf{w}, \mathbf{u}$ are dependent. Right: $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ lie in the plane $z = 0$ and are dependent; $\mathbf{u} = (0, 0, 1)$ leaves the plane, and $\mathbf{v}_1, \mathbf{v}_2, \mathbf{u}$ are independent.

(the added terms are $\mathbf{0}$ by Proposition 1.2.8(d) and rule (3)), a relation with the weight $-1 \neq 0$.

($\Leftarrow$) Again the contrapositive. Suppose $c_1\mathbf{v}_1 + \cdots + c_k\mathbf{v}_k = \mathbf{0}$ with not all c_i zero, and let j be the *largest* index with $c_j \neq 0$. The terms after $c_j\mathbf{v}_j$ are $0\mathbf{v}_i = \mathbf{0}$, so (rule (3)) $c_1\mathbf{v}_1 + \cdots + c_j\mathbf{v}_j = \mathbf{0}$. If $j = 1$, this says $c_1\mathbf{v}_1 = \mathbf{0}$ with $c_1 \neq 0$, so $\mathbf{v}_1 = \mathbf{0}$ by Proposition 1.2.8(g). If $j \geq 2$, the second half of the proof of Theorem 2.6.3, applied to $\mathbf{v}_1, \dots, \mathbf{v}_j$, gives $\mathbf{v}_j = (-c_1/c_j)\mathbf{v}_1 + \cdots + (-c_{j-1}/c_j)\mathbf{v}_{j-1} \in \text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_{j-1}\}$. Either way the condition of the theorem fails. $\square$

Combined with Theorem 1.3.11, this says that in an independent list every vector counts: each one enlarges the span of the ones before it, and leaving out any one of them makes the span smaller. (If leaving out $\mathbf{v}_j$ did not shrink the span, $\mathbf{v}_j$ would lie in the span of the others, and the list would be dependent by Theorem 2.6.3.)

2.6.5 Testing independence by elimination

For larger lists, pictures are not enough. Let A be the $n \times k$ matrix whose columns are $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{R}^n$, so a_{ij} is the i -th entry of $\mathbf{v}_j$. The i -th entry of $c_1\mathbf{v}_1 + \cdots + c_k\mathbf{v}_k$ is $a_{i1}c_1 + \cdots + a_{ik}c_k$, so the vector equation $c_1\mathbf{v}_1 + \cdots + c_k\mathbf{v}_k = \mathbf{0}$ is exactly the linear system

$$\left[A \mid \mathbf{0} \right]$$

in the unknowns $c_1, \dots, c_k$. A system whose right-hand side is $\mathbf{0}$ is *homogeneous* (Section 2.5); it always has the *trivial solution* $c_1 = \cdots = c_k = 0$. Row operations never change a column of zeros (Proposition 2.5.8: each new entry is $0, c \cdot 0$ or $0 + c \cdot 0$), so the last column can be left out and only A needs to be row reduced.

Theorem 2.6.8 (Elimination test). Let A be the matrix with columns $\mathbf{v}_1, \dots, \mathbf{v}_k$, and let E be a row echelon form of A . The following are equivalent:

- (a) $\mathbf{v}_1, \dots, \mathbf{v}_k$ are linearly independent;
- (b) the homogeneous system $[A \mid \mathbf{0}]$ has only the trivial solution;
- (c) every column of E is a pivot column.

If column j of E has no pivot, setting $c_j = 1$, every other non-pivot unknown to 0, and solving

for the pivot unknowns by back substitution gives a linear dependence relation.

Proof. (a) $\Leftrightarrow$ (b): by the computation above, the solutions of the system are exactly the lists of weights $(c_1, \dots, c_k)$ with $c_1 \mathbf{v}_1 + \dots + c_k \mathbf{v}_k = \mathbf{0}$.

(b) $\Leftrightarrow$ (c): this is Proposition 2.5.8 (a nontrivial solution exists exactly when there is a free variable); we give the argument in full, because it also produces the relation. By Theorem 2.3.4, $[A \mid \mathbf{0}]$ and $[E \mid \mathbf{0}]$ have the same solutions, so we work with E . Write $p_1 < p_2 < \dots < p_r$ for the pivot columns, where the pivot of row i is $e_{ip_i} \neq 0$ and rows $r+1, \dots, n$ are zero (they say $0 = 0$).

Suppose every column is a pivot column. Then $r = k$, and since $1 \leq p_1 < \dots < p_k \leq k$, we have $p_i = i$: the pivots are $e_{11}, \dots, e_{kk}$. Let $(c_1, \dots, c_k)$ be a solution. Row k says $e_{kk}c_k = 0$, so $c_k = 0$ (multiply by e_{kk}^{-1}). Going up, if $c_{i+1} = \dots = c_k = 0$, row i says $e_{ii}c_i + e_{i,i+1}c_{i+1} + \dots + e_{ik}c_k = e_{ii}c_i = 0$, so $c_i = 0$. Hence only the trivial solution exists.

Suppose column j has no pivot. Set $c_j = 1$ and every other non-pivot unknown to 0. Then go up through rows $r, r-1, \dots, 1$: row i has the form $e_{ip_i}c_{p_i} + (\text{terms in } c_l \text{ with } l > p_i) = 0$, and every c_l with $l > p_i$ is either a non-pivot unknown (already set) or the pivot unknown of a lower row (already found), so

$$c_{p_i} = -e_{ip_i}^{-1} (e_{i,p_i+1}c_{p_i+1} + \dots + e_{ik}c_k)$$

is determined. All r equations hold, so this is a solution, and it is not trivial because $c_j = 1$. This proves the last sentence of the theorem and, with it, (b) $\Rightarrow$ (c). $\square$

Example 2.6.9. Let $\mathbf{v}_1 = (1, 2, 3)$, $\mathbf{v}_2 = (4, 5, 6)$, $\mathbf{v}_3 = (7, 8, 9)$. Row reduce the matrix with these columns:

$$\begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix} \xrightarrow[\substack{R_2 \rightarrow R_2 - 2R_1 \\ R_3 \rightarrow R_3 - 3R_1}]{\substack{R_2 \rightarrow R_2 - 2R_1 \\ R_3 \rightarrow R_3 - 3R_1}} \begin{bmatrix} 1 & 4 & 7 \\ 0 & -3 & -6 \\ 0 & -6 & -12 \end{bmatrix} \xrightarrow{R_3 \rightarrow R_3 - 2R_2} \begin{bmatrix} 1 & 4 & 7 \\ 0 & -3 & -6 \\ 0 & 0 & 0 \end{bmatrix}.$$

(Row 2: $2 - 2 \cdot 1 = 0$, $5 - 2 \cdot 4 = -3$, $8 - 2 \cdot 7 = -6$. Row 3: $3 - 3 \cdot 1 = 0$, $6 - 3 \cdot 4 = -6$, $9 - 3 \cdot 7 = -12$, then $0 - 2 \cdot 0 = 0$, $-6 - 2(-3) = 0$, $-12 - 2(-6) = 0$.) The pivots are in columns 1 and 2; column 3 has none, so c_3 is free and the vectors are *dependent*. With $c_3 = 1$: row 2 says $-3c_2 - 6 = 0$, so $c_2 = -2$; row 1 says $c_1 + 4(-2) + 7 = c_1 - 1 = 0$, so $c_1 = 1$. Hence

$$\mathbf{v}_1 - 2\mathbf{v}_2 + \mathbf{v}_3 = \mathbf{0}, \quad \text{check: } (1 - 8 + 7, 2 - 10 + 8, 3 - 12 + 9) = (0, 0, 0).$$

A dependence relation is a column relation of A in the sense of Section 2.4, here $\mathbf{v}_3 = -\mathbf{v}_1 + 2\mathbf{v}_2$, and by Proposition 2.4.11 the echelon form keeps it: its third column is again minus the first plus twice the second, $(7, -6, 0) = -(1, 0, 0) + 2(4, -3, 0)$, since $-1 + 8 = 7$, $0 - 6 = -6$ and $0 + 0 = 0$. So the relation can be read off the echelon form directly.

Now replace $\mathbf{v}_3$ by $\mathbf{w} = (7, 8, 10)$, which differs only in its last entry. The same row operations change only the third column: $8 - 2 \cdot 7 = -6$, $10 - 3 \cdot 7 = -11$, and then $-11 - 2(-6) = 1$:

$$\begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 10 \end{bmatrix} \longrightarrow \begin{bmatrix} 1 & 4 & 7 \\ 0 & -3 & -6 \\ 0 & 0 & 1 \end{bmatrix}.$$

Every column has a pivot, so $\mathbf{v}_1, \mathbf{v}_2, \mathbf{w}$ are *independent*.

2.6.6 Too many vectors

Theorem 2.6.10. If $k > n$, then any k vectors in $\mathbb{F}^n$ are linearly dependent.

Proof. The matrix A with these columns is $n \times k$. In a row echelon form each pivot is the leading entry of a different row, so there are at most n pivots. Since $k > n$, at least one of the k columns has no pivot (Figure 2.10), and the vectors are dependent by Theorem 2.6.8. (Equivalently: the homogeneous system $[A \mid \mathbf{0}]$ has k unknowns and only $n < k$ equations, so it has a nontrivial solution by Theorem 2.5.9.) $\square$




Figure 2.10: Four vectors in $\mathbb{R}^3$: a row echelon form of the 3×4 matrix has at most three pivots, so some column has no pivot and the vectors are dependent (Theorem 2.6.10).

For example, any three vectors in $\mathbb{R}^2$ are dependent; for $(1, 0)$, $(0, 1)$, $(2, 3)$ the relation is $2(1, 0) + 3(0, 1) - (2, 3) = \mathbf{0}$, since $2 \cdot 1 + 3 \cdot 0 = 2$ and $2 \cdot 0 + 3 \cdot 1 = 3$. The converse is false: at most n vectors may well be dependent, as $(1, 2, 3)$, $(4, 5, 6)$, $(7, 8, 9)$ in Example 2.6.9 show. For exactly n vectors in $\mathbb{F}^n$ the matrix is square, and the vectors are independent exactly when an echelon form has n pivots, one in every row and every column. In Lecture 17 (Section 3.6) this condition joins a long list of equivalent ones in the invertible matrix theorem. And since $\mathbb{F}^n$ has room for at most n independent vectors, n will turn out to be its *dimension* (Lecture 20).

2.6.7 The field matters

Definition 2.6.1 uses only sums, scalar multiples and the zero vector, so it makes sense over every field, and the elimination test works word for word. But whether a list is independent can depend on the field.

Example 2.6.11. Let $\mathbf{a} = (1, 1, 0)$, $\mathbf{b} = (0, 1, 1)$, $\mathbf{c} = (1, 0, 1)$. Over $\mathbb{R}$, row reduce the matrix with these columns:

$$\begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \xrightarrow{R_2 \rightarrow R_2 - R_1} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & -1 \\ 0 & 1 & 1 \end{bmatrix} \xrightarrow{R_3 \rightarrow R_3 - R_2} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & -1 \\ 0 & 0 & 2 \end{bmatrix}$$

($1 - 1 = 0$, $1 - 0 = 1$, $0 - 1 = -1$; then $0 - 0 = 0$, $1 - 1 = 0$, $1 - (-1) = 2$). There are three pivots, so over $\mathbb{R}$ the vectors are independent. Over $\mathbb{Z}_2$, where $1 + 1 = 0$,

$$\mathbf{a} + \mathbf{b} + \mathbf{c} = (1 + 0 + 1, 1 + 1 + 0, 0 + 1 + 1) = (0, 0, 0),$$

a dependence relation with all weights 1; so over $\mathbb{Z}_2$ the same vectors are dependent. In the elimination over $\mathbb{Z}_2$, the last pivot 2 is 0 and column 3 has no pivot, in agreement with Theorem 2.6.8.

Check yourself

- (a) Are $(1, 1, 2)$, $(2, 3, 5)$, $(1, 2, 4)$ linearly independent? Decide by elimination.
 (b) Find a linear dependence relation among $(1, 2, 0)$, $(0, 1, 1)$, $(2, 5, 1)$.
 (c) True or false: if no two of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ are multiples of each other, then $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ are independent.

Supplement: beyond the lecture

Independence means unique weights. Let $S = \text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$. The vectors are independent if and only if every $\mathbf{x} \in S$ can be written as $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k$ in *only one* way. Indeed, if $\mathbf{x} = \sum c_i\mathbf{v}_i = \sum d_i\mathbf{v}_i$, subtracting gives $\sum (c_i - d_i)\mathbf{v}_i = \mathbf{0}$ (rules (1) to (4) and (8)), so independence forces $c_i = d_i$ for every i . Conversely, if the weights of $\mathbf{0}$ are unique, they must be the trivial ones, which is independence. This is why independent spanning lists, the *bases* of Lecture 20, give every vector well-defined coordinates (Section 5.1).

Nearly dependent vectors. On a computer, which rounds every result, the line between independent and dependent is blurred. The vectors $(1, 1)$ and $(1, 1 + 10^{-12})$ are independent, but only just: the replacement $R_2 \rightarrow R_2 - R_1$ leaves the pivot 10^{-12} , which is about the size of rounding errors in a longer calculation. Numerical software therefore asks not whether a pivot is zero but whether it is *small*, and measures how close a list is to dependence by its smallest singular value (Lecture 51).

Independence of functions. The definition makes sense whenever vectors can be added and scaled, for instance for functions (Lecture 21). The functions $1, x, x^2$ are independent: if $c_0 + c_1x + c_2x^2 = 0$ for every real x , then putting $x = 0$ gives $c_0 = 0$, and putting $x = 1$ and $x = -1$ gives $c_1 + c_2 = 0$ and $-c_1 + c_2 = 0$, so $c_1 = c_2 = 0$.

Answers to Check yourself**Section 2.1**

- (a) $3^{-1} = 5$ in $\mathbb{Z}_7$ ($3 \cdot 5 = 15 \rightarrow 1$). $3x + 2 = 6$ gives $3x = 4$, so $x = 5 \cdot 4 = 20 \rightarrow 6$. Check: $3 \cdot 6 + 2 = 20 \rightarrow 6$. (b) $(2 + i)^{-1} = \frac{2-i}{5} = \frac{2}{5} - \frac{1}{5}i$. Check: $(2 + i)(2 - i) = 4 + 1 = 5$, so $(2 + i) \cdot \frac{2-i}{5} = 1$. (c) No: $3 \cdot 3 = 9 \rightarrow 0$ in $\mathbb{Z}_9$, so 3 is a zero divisor and $\mathbb{Z}_9$ is not a field (Proposition 2.1.5(d); 9 is not prime, Theorem 2.1.16).

Section 2.2

- (a) $\left[\begin{array}{cc|c} 1 & 2 & 4 \\ 3 & -1 & 5 \end{array} \right]$. (b) $2 + 2 \cdot 1 = 4$ and $3 \cdot 2 - 1 = 5$. (c) $x(1, 3) + y(2, -1) = (4, 5)$, and $2(1, 3) + 1(2, -1) = (2, 6) + (2, -1) = (4, 5)$.

Section 2.3

- (a) $R_2 - R_1$ gives $[0 \ 1 \ 2 \mid 8]$, $R_3 - 2R_1$ gives $[0 \ 1 \ 3 \mid 11]$, $R_3 - R_2$ gives $[0 \ 0 \ 1 \mid 3]$; so $z = 3$, $y = 2$, $x = 1$. Check: $1 + 2 + 3 = 6$, $1 + 4 + 9 = 14$, $2 + 6 + 15 = 23$. (b) No: row 1's leading entry is to the right of row 2's. Swapping $R_1 \leftrightarrow R_2$ fixes it.

Section 2.4

- (a) The reduced row echelon form is $\left[\begin{array}{ccc|c} 1 & 2 & 0 & 4 \\ 0 & 0 & 1 & 1 \end{array} \right]$: x and z are basic, $y = t$ is free, so $x = 4 - 2t$, $y = t$, $z = 1$. Check: $t = 0$ gives $(4, 0, 1)$. (b) True: row equivalence is symmetric and transitive.

Section 2.5

- (a) Consistent, because the last column is not a pivot column. x_3 is free, so there are infinitely many

solutions: $\mathbf{x} = (2, 5, 0) + t(-3, 1, 1)$. The homogeneous solution set is $\text{Span}\{(-3, 1, 1)\}$. (b) There are at most 3 pivots among 5 columns, so at least 2 free variables (Proposition 2.5.5); a homogeneous system is always consistent, so it has a nontrivial solution (Theorem 2.5.9). A 3×5 system need not have a solution: $x_1 + \cdots + x_5 = 0$, $x_1 + \cdots + x_5 = 1$, $x_1 = 0$ is inconsistent (subtracting the first two equations gives $0 = 1$).

Section 2.6

- (a) Independent. The matrix with columns $(1, 1, 2)$, $(2, 3, 5)$, $(1, 2, 4)$ is $\begin{bmatrix} 1 & 2 & 1 \\ 1 & 3 & 2 \\ 2 & 5 & 4 \end{bmatrix}$; $R_2 - R_1$ gives $(0, 1, 1)$, $R_3 - 2R_1$ gives $(0, 1, 2)$, then $R_3 - R_2$ gives $(0, 0, 1)$. Three pivots, so the columns are independent (Theorem 2.6.8). (b) $2(1, 2, 0) + (0, 1, 1) - (2, 5, 1) = \mathbf{0}$. Check: $2 + 0 - 2 = 0$, $4 + 1 - 5 = 0$, $0 + 1 - 1 = 0$. (c) False: $(1, 0)$, $(0, 1)$, $(1, 1)$ contain no two multiples of each other, yet $(1, 0) + (0, 1) - (1, 1) = \mathbf{0}$.