

Dot Product and Cross Product

[Read this note online](#)

Dot product and cross product are the two fundamental ways to multiply vectors. They're not just calculation rules, they're carriers of geometric meaning. The dot product encodes angle and projection; the cross product encodes perpendicular direction and area. Together they handle nearly every vector calculation in physics, computer graphics, and engineering.


The dot product produces a scalar from two vectors. The cross product (defined only in 3D) produces a vector perpendicular to both inputs. Each captures a different aspect of how vectors relate, and learning to switch between them quickly is one of the practical superpowers of linear algebra.

The dot product measures alignment through a scalar. In 3-dimensional Euclidean space, the cross product produces a perpendicular vector whose orientation depends on the order.

DOT PRODUCT

$$\vec{u} \cdot \vec{v} = |\vec{u}| |\vec{v}| \cos(\theta)$$

$$= u_1 v_1 + u_2 v_2 + u_3 v_3$$



WHERE:

- $\vec{u} = \langle u_1, u_2, u_3 \rangle$
- $\vec{v} = \langle v_1, v_2, v_3 \rangle$
- θ = angle between $\vec{u}$ and $\vec{v}$

SPECIAL CASES:

- $\theta = 0^\circ$ (same direction)
 $\vec{u} \cdot \vec{v} = |\vec{u}| |\vec{v}|$
- $\theta = 90^\circ$ (perpendicular)
 $\vec{u} \cdot \vec{v} = 0$
- $\theta = 180^\circ$ (opposite direction)
 $\vec{u} \cdot \vec{v} = -|\vec{u}| |\vec{v}|$

PROPERTIES:

- $\vec{u} \cdot \vec{v} = \vec{v} \cdot \vec{u}$ (commutative)
- $\vec{u} \cdot (\vec{v} + \vec{w}) = \vec{u} \cdot \vec{v} + \vec{u} \cdot \vec{w}$ (distributive)
- $\vec{u} \cdot \vec{u} = |\vec{u}|^2$
- $\vec{u} \cdot \vec{v} = 0$ if and only if $\vec{u} \perp \vec{v}$




GEOMETRIC INTERPRETATION:

$\vec{u} \cdot \vec{v} = (\text{magnitude of } \vec{u}) (\text{projection of } \vec{v} \text{ onto } \vec{u})$

CROSS PRODUCT

$$\vec{u} \times \vec{v} = |\vec{u}| |\vec{v}| \sin(\theta) \hat{n}$$



WHERE:

- θ = angle from $\vec{u}$ to $\vec{v}$ (right-hand sense)
- $\hat{n}$ = unit vector perpendicular to both $\vec{u}$ and $\vec{v}$

PROPERTIES:

- $\vec{u} \times \vec{v} = -\vec{v} \times \vec{u}$ (anti-commutative)
- $\vec{u} \times (\vec{v} + \vec{w}) = \vec{u} \times \vec{v} + \vec{u} \times \vec{w}$
- $\vec{u} \times \vec{u} = \vec{0}$
- $\vec{u} \times \vec{v} = \vec{0}$ if and only if $\vec{u}$ and $\vec{v}$ are parallel





GEOMETRIC INTERPRETATION:

$|\vec{u} \times \vec{v}| = \text{area of the parallelogram spanned by } \vec{u} \text{ and } \vec{v}$

RIGHT-HAND RULE

1. Point your thumb in the direction of $\vec{u}$.
2. Curl your fingers toward $\vec{v}$ (through the smaller angle θ).
3. Your thumb points in the direction of $\vec{u} \times \vec{v}$ (or $\hat{n}$).



Dot product and cross product formulas right-hand rule modern textbook illustration

The Dot Product Formula

The dot product (inner product, scalar product) of two **vectors** is defined two equivalent ways:

The first form (component-wise sum) is computational. The second form (magnitudes times cosine of angle) is geometric. Both produce the same scalar.

The component formula generalizes to any number of dimensions. The geometric formula assumes the standard Euclidean structure but extends to inner-product spaces in functional analysis.

$$\mathbf{u} \cdot \mathbf{v} = u_1 v_1 + u_2 v_2 + \dots + u_n v_n = |\mathbf{u}| |\mathbf{v}| \cos \theta$$

Geometric Meaning of the Dot Product

The dot product captures how much one vector lies in the direction of the other. If $\mathbf{u} \cdot \mathbf{v} > 0$, the vectors point in roughly the same direction (angle less than 90°). If $\mathbf{u} \cdot \mathbf{v} < 0$, they point in roughly opposite directions (angle greater than 90°). If $\mathbf{u} \cdot \mathbf{v} = 0$, they're perpendicular.

Geometrically, $\mathbf{u} \cdot \mathbf{v}$ equals $|\mathbf{u}|$ times the length of the projection of $\mathbf{v}$ onto $\mathbf{u}$ (or vice versa). This makes the dot product a measure of "alignment", how parallel the vectors are.

For unit vectors, the dot product reduces to just $\cos \theta$, making the dot product the standard tool for computing angles between vectors.

Properties of the Dot Product

The Cauchy-Schwarz inequality is one of the most useful inequalities in mathematics; it underpins many proofs in analysis, statistics, and probability theory.

Commutative: $\mathbf{u} \cdot \mathbf{v} = \mathbf{v} \cdot \mathbf{u}$. **Distributive over addition:** $\mathbf{u} \cdot (\mathbf{v} + \mathbf{w}) = \mathbf{u} \cdot \mathbf{v} + \mathbf{u} \cdot \mathbf{w}$. **Compatible with scalar multiplication:** $(c\mathbf{u}) \cdot \mathbf{v} = c(\mathbf{u} \cdot \mathbf{v}) = \mathbf{u} \cdot (c\mathbf{v})$. **Self-dot equals magnitude squared:** $\mathbf{v} \cdot \mathbf{v} = |\mathbf{v}|^2$. **Cauchy-Schwarz inequality:** $|\mathbf{u} \cdot \mathbf{v}| \leq |\mathbf{u}| |\mathbf{v}|$, with equality if and only if the vectors are parallel.

Worked Example: Dot Product

Compute $\mathbf{u} \cdot \mathbf{v}$ and the angle between them, where $\mathbf{u} = (3, 4, 0)$ and $\mathbf{v} = (1, 0, 2)$.

Component-wise:

Magnitudes: $|\mathbf{u}| = \sqrt{9 + 16} = 5$, $|\mathbf{v}| = \sqrt{1 + 4} = \sqrt{5}$.

Angle: $\cos \theta = 3/(5 \cdot \sqrt{5}) \approx 0.268$, so $\theta \approx 74.5^\circ$.

The dot product gives both a quick orthogonality check (positive here, so the vectors aren't perpendicular) and a clean angle calculation when normalized.

$$\mathbf{u} \cdot \mathbf{v} = 3 \cdot 1 + 4 \cdot 0 + 0 \cdot 2 = 3$$

The Cross Product Formula

The cross product is defined only in 3D (and 7D for separate algebraic reasons). For $\mathbf{u}, \mathbf{v} \in \mathbb{R}^3$:

This produces a new vector perpendicular to both inputs. The magnitude is:

which equals the area of the parallelogram spanned by $\mathbf{u}$ and $\mathbf{v}$. Direction follows the right-hand rule.

$$\mathbf{u} \times \mathbf{v} = (u_2v_3 - u_3v_2, u_3v_1 - u_1v_3, u_1v_2 - u_2v_1)$$

$$|\mathbf{u} \times \mathbf{v}| = |\mathbf{u}||\mathbf{v}| \sin \theta$$

The Right-Hand Rule

To find the direction of $\mathbf{u} \times \mathbf{v}$, point your right hand's fingers along $\mathbf{u}$ and curl them toward $\mathbf{v}$. Your thumb points in the direction of the cross product.

Equivalent statement: $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{u} \times \mathbf{v}$ form a right-handed coordinate system, with the cross product playing the role of the third axis.

The right-hand rule is convention. Some physics communities (especially older European ones) use a left-hand rule for certain quantities; the math is the same up to sign. Modern textbooks universally adopt the right-hand convention.

Properties of the Cross Product

The non-associativity is a classic gotcha. It's why parenthesization matters in expressions involving multiple cross products and why the Jacobi identity (the cyclic relation

$$\mathbf{u} \times (\mathbf{v} \times \mathbf{w}) + \mathbf{v} \times (\mathbf{w} \times \mathbf{u}) + \mathbf{w} \times (\mathbf{u} \times \mathbf{v}) = \mathbf{0}$$

) is the formal compensation.

Anticommutative: $\mathbf{u} \times \mathbf{v} = -(\mathbf{v} \times \mathbf{u})$. Order matters and reversing it flips the sign. **Distributive:** $\mathbf{u} \times (\mathbf{v} + \mathbf{w}) = \mathbf{u} \times \mathbf{v} + \mathbf{u} \times \mathbf{w}$. **Scalar multiplication:** $(c\mathbf{u}) \times \mathbf{v} = c(\mathbf{u} \times \mathbf{v})$. **Not associative:**

$(\mathbf{u} \times \mathbf{v}) \times \mathbf{w} \neq \mathbf{u} \times (\mathbf{v} \times \mathbf{w})$ in general. **Cross product with itself is zero:** $\mathbf{v} \times \mathbf{v} = \mathbf{0}$. **Cross product of parallel vectors is zero:** if $\mathbf{u}$ and $\mathbf{v}$ are parallel, $\mathbf{u} \times \mathbf{v} = \mathbf{0}$.

Worked Example: Cross Product

Compute $\mathbf{u} \times \mathbf{v}$ for $\mathbf{u} = (1, 2, 3)$ and $\mathbf{v} = (4, 5, 6)$.

Using the formula:

Verify perpendicularity:

$$\mathbf{u} \cdot (\mathbf{u} \times \mathbf{v}) = 1 \cdot (-3) + 2 \cdot 6 + 3 \cdot (-3) = -3 + 12 - 9 = 0.$$

$$\mathbf{v} \cdot (\mathbf{u} \times \mathbf{v}) = 4 \cdot (-3) + 5 \cdot 6 + 6 \cdot (-3) = -12 + 30 - 18 = 0.$$

The result is perpendicular to both inputs, as required.

$$\mathbf{u} \times \mathbf{v} = (2 \cdot 6 - 3 \cdot 5, 3 \cdot 4 - 1 \cdot 6, 1 \cdot 5 - 2 \cdot 4) = (-3, 6, -3)$$

Determinant Form for the Cross Product

The cross product can be written as a **determinant**:

This is a mnemonic, not a formal determinant (the first row contains vectors, not scalars). Expanding along the first row reproduces the standard cross product formula.

The determinant form makes some properties easier to see. Two parallel vectors give two identical rows, hence zero determinant, explaining why parallel vectors have zero cross product. Switching $\mathbf{u}$ and $\mathbf{v}$ swaps the second and third rows, flipping the sign, explaining anticommutativity.

$$\mathbf{u} \times \mathbf{v} = \det \begin{pmatrix} \hat{i} & \hat{j} & \hat{k} \\ u_1 & u_2 & u_3 \\ v_1 & v_2 & v_3 \end{pmatrix}$$

The Triple Product

The scalar triple product $\mathbf{u} \cdot (\mathbf{v} \times \mathbf{w})$ gives the signed volume of the parallelepiped spanned by the three vectors. It equals the determinant:

The triple product is invariant under cyclic permutation: $\mathbf{u} \cdot (\mathbf{v} \times \mathbf{w}) = \mathbf{v} \cdot (\mathbf{w} \times \mathbf{u}) = \mathbf{w} \cdot (\mathbf{u} \times \mathbf{v})$. It changes sign under non-cyclic swaps.

The vector triple product $\mathbf{u} \times (\mathbf{v} \times \mathbf{w})$ satisfies the BAC-CAB rule: $\mathbf{u} \times (\mathbf{v} \times \mathbf{w}) = \mathbf{v}(\mathbf{u} \cdot \mathbf{w}) - \mathbf{w}(\mathbf{u} \cdot \mathbf{v})$.

This is one of the most useful identities in vector calculus.

$$\mathbf{u} \cdot (\mathbf{v} \times \mathbf{w}) = \det \begin{pmatrix} u_1 & u_2 & u_3 \\ v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \end{pmatrix}$$

Why the Cross Product Only Exists in 3D

In 2D, "perpendicular to two vectors" doesn't define a unique direction within the plane, the perpendicular points out of the plane. In 4D and higher, "perpendicular to two given vectors" is a multidimensional subspace, not a single direction. So a single perpendicular vector valued operation doesn't exist there.

The exception is 7D, where a related cross product exists due to the algebraic structure of octonions. This is largely a curiosity of higher mathematics; for everyday work, the cross product is a 3D-only operation.

In 2D, the analogue is the "perpendicular dot product" (also called the wedge product or 2D cross product), which gives a scalar equal to the signed area of the parallelogram. It's not a vector, but it captures the same area information that the 3D cross product magnitude captures.

Where Dot and Cross Products Show Up

Physics: work = force · displacement (dot product); torque = $\mathbf{r} \times \mathbf{F}$ (cross product); angular momentum $\mathbf{L} = \mathbf{r} \times \mathbf{p}$ (cross product); magnetic force $\mathbf{F} = q\mathbf{v} \times \mathbf{B}$ (cross product). **Computer graphics:** lighting calculations (Lambertian shading uses light direction · normal); backface culling (sign of normal · view direction); surface normal computation (cross product of two edge vectors). **Machine learning:** cosine similarity uses dot products to measure similarity between vectors (recommendation systems, document retrieval, embedding alignment). **Engineering:** work-energy calculations; structural analysis; fluid mechanics where vorticity involves cross products. **Geometry:** determining whether points are on the same side of a line (sign of cross product); area and volume calculations; orientation tests. **Robotics:** joint angle computations and end-effector kinematics rely heavily on dot and cross products.

Cosine Similarity in Machine Learning

Cosine similarity measures how aligned two vectors are, ignoring their magnitudes:

It returns a value between -1 and 1. Cosine similarity is the standard metric for comparing word embeddings, document vectors, image feature vectors, and recommendation systems where di-

rection matters more than scale.

Why magnitude doesn't matter: in many embedding systems, vector magnitude reflects something other than semantic content (e.g., document length). Stripping out magnitude lets the comparison focus on what the vectors mean rather than how large they are. Read more in [best machine learning courses](#).

$$\cos_sim(\mathbf{u}, \mathbf{v}) = \frac{\mathbf{u} \cdot \mathbf{v}}{|\mathbf{u}||\mathbf{v}|} = \cos \theta$$

Common Mistakes with Dot and Cross Products

Confusing the two operations. Dot product produces a scalar; cross product produces a vector. They have different geometric meanings. **Forgetting that the cross product is anticommutative.** $\mathbf{u} \times \mathbf{v} = -(\mathbf{v} \times \mathbf{u})$. The order matters. **Trying to compute a 2D cross product as a vector.** The 2D cross product is a scalar (area), not a vector. The vector cross product needs at least 3 dimensions. **Assuming cross product is associative.** It isn't. Use the BAC-CAB rule for vector triple products. **Misapplying the right-hand rule.** The fingers point along the first vector and curl toward the second. Reversing this gives the wrong direction. **Forgetting that orthogonal vectors have zero dot product, not zero cross product.** Perpendicular vectors have maximum cross product magnitude (they form a parallelogram with the largest possible area).

A Brief History

The dot product traces back to William Rowan Hamilton's quaternion algebra (1843) and Hermann Grassmann's exterior algebra (1844). Josiah Willard Gibbs and Oliver Heaviside extracted the modern dot and cross products from these systems in the 1880s, presenting them in the form we use today.

Gibbs' *Vector Analysis* (1881) and Heaviside's *Electromagnetic Theory* (1893) made vector notation standard in physics and engineering. The modern format, dot product as scalar, cross product as vector, has been the textbook standard for over a century.

Inner Products in General Vector Spaces

The dot product generalizes to inner products on abstract vector spaces. An inner product is any positive-definite, symmetric, bilinear form. Function spaces have inner products like $\langle f, g \rangle = \int f(x)g(x) dx$. Polynomial spaces, matrix spaces, and infinite-dimensional Hilbert spaces all support inner-product geometry.

Inner products give length ($\|f\| = \sqrt{\langle f, f \rangle}$) and angle ($\cos \theta = \langle f, g \rangle / (\|f\| \|g\|)$) in any vector space.

This is the foundation of Hilbert space theory, Fourier analysis, and quantum mechanics. The familiar Euclidean dot product is just the simplest example of a much broader concept.

The Wedge Product and Its Generalizations

The wedge product (exterior product) generalizes the cross product to any dimension. $\mathbf{u} \wedge \mathbf{v}$ is a "bivector", an oriented area element, rather than a vector. In 3D, the wedge product is dual to the cross product (a single perpendicular vector encodes the bivector).

Wedge products extend cleanly to higher dimensions and produce the right geometric objects (areas in 2D, volumes in 3D, hypervolumes in n D). They underlie differential forms, exterior calculus, and the most general formulations of multivariable calculus and physics.

Gram-Schmidt Orthogonalization

The Gram-Schmidt process turns any linearly independent set of vectors into an orthonormal set spanning the same subspace. Each step uses dot products to subtract out the components already covered by earlier vectors, leaving a vector orthogonal to all previous ones, which is then normalized to unit length.

Gram-Schmidt is the foundation of QR decomposition and many least-squares regression algorithms. It's also a clean illustration of dot products in action, every step of the algorithm computes dot products, projections, and subtractions to enforce orthogonality. Numerical implementations use modified Gram-Schmidt or Householder reflections for stability, but the underlying idea is the same.

Cosine Similarity in Practice

Cosine similarity computes $\text{cos_sim}(\mathbf{u}, \mathbf{v}) = \mathbf{u} \cdot \mathbf{v} / (|\mathbf{u}| |\mathbf{v}|)$. It returns values in $[-1, 1]$: 1 means identical direction, 0 means perpendicular, -1 means opposite.

Cosine similarity is the standard tool for comparing high-dimensional embeddings, word vectors (word2vec, GloVe), sentence embeddings (BERT, GPT), image embeddings (CLIP), and user/item vectors in recommender systems. It's preferred over Euclidean distance because it ignores magnitude, which often reflects irrelevant factors (document length, signal strength).

Computing cosine similarity at scale is well-supported by vector databases (Pinecone, Weaviate, Milvus, FAISS). Approximate nearest-neighbor algorithms speed up cosine-similarity search across billions of vectors.

Worked Example: Torque Calculation

A wrench applies force $\mathbf{F} = (10, 0, 0)$ Newtons at the end of a lever arm $\mathbf{r} = (0, 0.3, 0)$ meters from a bolt. Compute the torque.

Torque $\tau = \mathbf{r} \times \mathbf{F}$:

The torque is 3 N·m about the negative z-axis. The magnitude is 3 (consistent with $|\mathbf{r}||\mathbf{F}| \sin 90^\circ = 0.3 \cdot 10$). The direction (negative z) tells you the bolt rotates clockwise as seen from above.

Every torque calculation in mechanical engineering uses the cross product this way. The same math underlies angular momentum, magnetic force, and rotational dynamics.

$$\tau = (0.3 \cdot 0 - 0 \cdot 0, 0 \cdot 10 - 0 \cdot 0, 0 \cdot 0 - 0.3 \cdot 10) = (0, 0, -3) \text{ N}\cdot\text{m}$$

The related [OpenStax mathematics textbooks](#) provide additional explanations and exercises for the mathematical definitions used here.

A Projection and an Area

For $a = (2, 1, 0)$ and $b = (1, 3, 0)$, the dot product is 5. The component of a along the unit direction of b is $5/\sqrt{10}$. The vector projection is $(a \cdot b)/(b \cdot b) b = (1/2, 3/2, 0)$.

The cross product is $a \times b = (0, 0, 5)$. Its magnitude is the area of the parallelogram spanned by the vectors. Reversing the order gives $(0, 0, -5)$, so the area stays the same while the oriented normal reverses.

The points to check in this example are:

- A scalar projection and a vector projection are different outputs.
- Perpendicular vectors have zero dot product.
- Parallel nonzero vectors have zero cross product.

Check Your Understanding

Try the questions before reading the answers.

Question 1. What is $a \cdot a$?

Answer. 5, the squared magnitude of a.

Question 2. What is $b \times a$?

Answer. $(0, 0, -5)$.

Practice Questions

Work each question before reading its solution. The set runs from direct recall and substitution to the applied questions that exams actually use to separate grades. The downloadable PDF includes these questions and their solutions.

Question 1. Compute $\vec{a} \cdot \vec{b}$ for $\vec{a} = (2, 3, -1)$ and $\vec{b} = (1, -2, 4)$.

Solution. Multiply matching components and add: $2(1) + 3(-2) + (-1)(4) = 2 - 6 - 4 = -8$. The negative sign already tells you the angle between the vectors is obtuse.

Question 2. Find the angle between $\vec{u} = (1, 0)$ and $\vec{v} = (1, 1)$.

Solution. $\cos \theta = \frac{\vec{u} \cdot \vec{v}}{|\vec{u}||\vec{v}|} = \frac{1}{1 \cdot \sqrt{2}} = \frac{1}{\sqrt{2}}$, so $\theta = 45^\circ$. The dot product is the standard tool for extracting angles.

Question 3. Are $\vec{a} = (3, -2, 1)$ and $\vec{b} = (2, 3, 0)$ perpendicular?

Solution. $\vec{a} \cdot \vec{b} = 6 - 6 + 0 = 0$. Yes. A zero dot product between nonzero vectors means exactly perpendicular, and this check is faster than computing any angle.

Question 4. Find the projection of $\vec{a} = (3, 4)$ onto $\vec{b} = (1, 0)$.

Solution. $\text{proj}_{\vec{b}} \vec{a} = \frac{\vec{a} \cdot \vec{b}}{|\vec{b}|^2} \vec{b} = \frac{3}{1}(1, 0) = (3, 0)$. Projection answers "how much of $\vec{a}$ points along $\vec{b}$ ", here simply the x-component because $\vec{b}$ is the x-axis direction.

Question 5. Compute $\hat{i} \times \hat{j}$ and state the general rule it illustrates.

Solution. $\hat{i} \times \hat{j} = \hat{k}$. The cross product of 2 of the coordinate unit vectors gives the third, with sign following the right-hand rule and the cyclic order $i \rightarrow j \rightarrow k$. Reversing the order flips the sign: $\hat{j} \times \hat{i} = -\hat{k}$.

Question 6. For $\vec{a} = (2, 1, 0)$ and $\vec{b} = (1, 3, 0)$, find $\vec{a} \times \vec{b}$ and the area of the parallelogram they span.

Solution.

$$\vec{a} \times \vec{b} = (1 \cdot 0 - 0 \cdot 3, 0 \cdot 1 - 2 \cdot 0, 2 \cdot 3 - 1 \cdot 1) = (0, 0, 5)$$

The magnitude 5 is the parallelogram area; the triangle on the same 2 vectors has half that, 2.5. The direction $(0, 0, 5)$ is perpendicular to both inputs, as a cross product must be.

Question 7. A force $\vec{F} = (4, 2)$ N moves an object through displacement $\vec{d} = (3, 1)$ m. Find the work done.

Solution. $W = \vec{F} \cdot \vec{d} = 12 + 2 = 14$ J. Work is the textbook dot-product application: only the force component along the motion counts.

Question 8. A wrench handle of length 0.4 m is pulled with a 25 N force at right angles to the

handle. Find the torque magnitude.

Solution. $|\vec{\tau}| = |\vec{r}||\vec{F}| \sin \theta = 0.4 \times 25 \times \sin 90^\circ = 10 \text{ N}\cdot\text{m}$. Torque is the textbook cross-product application: the perpendicular component of force does the turning, which is why pulling along the handle ($\sin 0^\circ = 0$) turns nothing.

Question 9. Find the volume of the parallelepiped spanned by $\vec{a} = (1, 0, 0)$, $\vec{b} = (0, 2, 0)$, $\vec{c} = (0, 0, 3)$.

Solution. Volume = $|\vec{a} \cdot (\vec{b} \times \vec{c})|$. Here $\vec{b} \times \vec{c} = (6, 0, 0)$, and $\vec{a} \cdot (6, 0, 0) = 6$. The scalar triple product measures signed volume, and these 3 axis-aligned edges form a 1 by 2 by 3 box.

Question 10. Why does the cross product exist only in 3 dimensions, while the dot product works in any dimension?

Solution. The dot product only needs matching components to multiply and add, which works in any dimension. The cross product must return a vector perpendicular to both inputs, and only in 3D is that perpendicular direction unique up to sign. In 2D there is no third direction; in 4D there are infinitely many.

FAQs

What is the difference between dot product and cross product?

Dot product takes two vectors and returns a scalar. Cross product (defined only in 3D) takes two vectors and returns a vector perpendicular to both. Dot product captures alignment and angle; cross product captures area and perpendicular direction.

How do I compute the dot product?

Sum the products of corresponding components: $\mathbf{u} \cdot \mathbf{v} = u_1v_1 + u_2v_2 + \dots + u_nv_n$. Equivalently, it equals $|\mathbf{u}||\mathbf{v}|\cos(\theta)$, where θ is the angle between the vectors. The component formula is the standard way to compute it.

How do I compute the cross product?

Use the formula $\mathbf{u} \times \mathbf{v} = (u_2v_3 - u_3v_2, u_3v_1 - u_1v_3, u_1v_2 - u_2v_1)$ or the determinant mnemonic with $\hat{i}, \hat{j}, \hat{k}$ in the first row. The result is a vector perpendicular to both $\mathbf{u}$ and $\mathbf{v}$ with magnitude $|\mathbf{u}||\mathbf{v}|\sin(\theta)$.

What does it mean if the dot product is zero?

The two vectors are perpendicular (orthogonal), assuming neither is zero. The dot product is the standard test for orthogonality, used everywhere from Gram-Schmidt orthogonalization to neural network weight initialization.

What does it mean if the cross product is zero?

The two vectors are parallel (or one is the zero vector). The cross product magnitude equals the area of the parallelogram spanned by the vectors; parallel vectors form a degenerate parallelogram with zero area.

Why does the cross product only exist in 3D?

Because 'perpendicular to two given vectors' uniquely defines a single direction (up to sign) only in 3D. In 2D, perpendicular points out of the plane. In higher dimensions, the perpendicular subspace is multidimensional. The exception is 7D, where a related cross product exists for algebraic reasons.

What is the right-hand rule?

Point your right hand's fingers along the first vector and curl them toward the second; your thumb points in the direction of the cross product. It's the convention for orienting cross products in 3D and ensures $u, v, u \times v$ form a right-handed coordinate system.

How is the dot product used in machine learning?

Cosine similarity (dot product divided by product of magnitudes) is the standard metric for comparing embeddings, recommendation matching, document similarity, and many other ML tasks. Neural network weights and inputs are combined via dot products at every layer.

How is the cross product used in physics?

Torque = $r \times F$. Angular momentum = $r \times p$. Magnetic force = $qv \times B$. Anywhere rotational or perpendicular geometry matters, cross products show up. In rigid body dynamics, the inertia tensor and angular velocity combine via cross products to give torque equations.

What is cosine similarity?

It's the cosine of the angle between two vectors: $\text{cos_sim}(u, v) = (u \cdot v) / (|u| \cdot |v|)$. It returns a value between -1 (opposite directions) and 1 (same direction), with 0 indicating perpendicular. Used everywhere from search engines to recommendation systems to embedding space analysis.

Is the dot product commutative?

Yes. $u \cdot v = v \cdot u$. The dot product is symmetric in its two arguments. The cross product, by contrast, is anticommutative: $u \times v = -(v \times u)$.

Can the dot product be negative?

Yes. Negative dot product means the vectors form an obtuse angle (greater than 90°). It happens naturally when vectors point in roughly opposite directions. Cosine similarity returns negative values in this case.

What's the geometric meaning of the dot product?

It equals $|u||v|\cos(\theta)$. For unit vectors, it's just the cosine of the angle between them. Geometrically, it captures how aligned the vectors are. Positive means roughly aligned; negative means roughly opposite; zero means perpendicular.

What's the difference between cross product and outer product?

Cross product takes two vectors and returns a third vector (in 3D). Outer product takes two vectors and returns a matrix: $(u v^T)_{ij} = u_i v_j$. The outer product is sometimes called the tensor product or Kronecker product in different contexts.