Module 1 — Hilbert Spaces

Quantum Foundations

hbar.university

Complex Vector Spaces

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Hilbert Spaces **Node:** 1.1

This is the foundation node of the entire course. Every subsequent structure — inner products, operators, observables, states, measurements, entanglement — is built on top of complex vector spaces. Real vector spaces are insufficient for quantum mechanics; the complex field is not optional.

Formal Definition

A **complex vector space** is a set V equipped with two operations:

1. Vector addition: $+ : V \times V \to V$
2. Scalar multiplication: $\cdot : \mathbb{C} \times V \to V$

satisfying, for all $u, v, w \in V$ and $\alpha, \beta \in \mathbb{C}$:

- **Associativity:** $(u + v) + w = u + (v + w)$
- **Commutativity:** $u + v = v + u$
- **Zero element:** there exists $0 \in V$ such that $v + 0 = v$
- **Additive inverse:** for each $v \in V$ there exists $-v \in V$ with $v + (-v) = 0$
- **Distributivity (vector):** $\alpha(u + v) = \alpha u + \alpha v$
- **Distributivity (scalar):** $(\alpha + \beta)v = \alpha v + \beta v$
- **Compatibility:** $\alpha(\beta v) = (\alpha\beta)v$
- **Identity:** $1 \cdot v = v$

A **basis** of V is a linearly independent set $\{e_i\} \subset V$ whose linear span equals V. The cardinality of any basis is the **dimension** of V, denoted $\dim V$.

Intuition

A complex vector space is a structure in which:

- Things can be added.
- Things can be scaled by complex numbers.
- The result is again a thing of the same kind.

The complex field is essential because quantum mechanical amplitudes interfere — and interference requires both magnitude and phase. A real vector space carries no notion of phase. This is the structural reason quantum mechanics cannot be formulated over $\mathbb{R}$.

Structural Role

Complex vector spaces are the carrier of quantum states. The state postulate (Node 2.1) will identify physical states with rays in a complex vector space, but the linear-algebraic structure is logically prior.

Every structural object in this course — operators, observables, measurements, channels, entropies — is ultimately a function of, or operator on, a complex vector space.

Mathematical Development

Linear combinations

For $v_1, \ldots, v_n \in V$ and $\alpha_1, \ldots, \alpha_n \in \mathbb{C}$, the **linear combination** is:

$$\sum_{i=1}^{n} \alpha_i v_i \in V.$$

Linear independence

A set $\{v_1, \ldots, v_n\}$ is **linearly independent** if

$$\sum_{i=1}^{n} \alpha_i v_i = 0 \implies \alpha_1 = \cdots = \alpha_n = 0.$$

Span and dimension

The **span** of a set $S \subset V$ is the set of all finite linear combinations of its elements. V is **finite-dimensional** if it has a finite basis; otherwise it is **infinite-dimensional**.

The standard example

$\mathbb{C}^n$ — the set of column vectors with n complex entries — is the prototype finite-dimensional complex vector space. The standard basis is $\{e_1, \ldots, e_n\}$ with e_i having a 1 in position i and 0 elsewhere.

Worked Example

The qubit space

The simplest non-trivial quantum system uses $V = \mathbb{C}^2$. A general element is

$$v = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \qquad \alpha, \beta \in \mathbb{C}.$$

The standard basis is

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

A general qubit vector is the linear combination

$$v = \alpha|0\rangle + \beta|1\rangle.$$

This already exhibits the central feature of quantum mechanics: a system can exist in a superposition of basis states with arbitrary complex coefficients.

Cross-Links

- **Linear Algebra:** Vector Spaces (Module 1) — same axiom system, real-or-complex field.
- **Statistics:** Probability simplex is a real convex set; quantum state space replaces it with a complex projective space.
- **VQA:** Parameterized circuits act on $\mathbb{C}^{2^n}$; the underlying linearity is what makes gradient-based optimization tractable.
- **Quantum Information:** Qubits, qudits, and quantum registers are all built directly on this object.

Structural Compression

Invariant: Closure under complex linear combination.

Mechanism: Two operations (addition, complex scalar multiplication) plus eight axioms.

Cross-System Echo: The same axioms over $\mathbb{R}$ give classical state spaces. Replacing $\mathbb{R}$ with $\mathbb{C}$ is the minimal structural perturbation that produces quantum behavior.

Exercises

1. Show that $\mathbb{C}^n$ satisfies all eight axioms of a complex vector space.
2. Prove that the zero element of a vector space is unique.
3. Show that the set of polynomials of degree $\leq n$ with complex coefficients forms a complex vector space, and identify its dimension.
4. Show that $\mathbb{C}$ is itself a complex vector space, and a real vector space — and that its dimension differs in the two cases.
5. Argue, structurally, why a real vector space cannot support the phenomenon of interference.

Inner Products and Norms

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Hilbert Spaces **Node:** 1.2

Inner products turn a complex vector space into a geometric object. Length, angle, and orthogonality become well-defined. The Born rule (Node 2.3) is fundamentally a statement about inner products — every probability in quantum mechanics is the squared modulus of one.

Formal Definition

A **complex inner product** on a complex vector space V is a map

$$\langle \cdot, \cdot \rangle : V \times V \to \mathbb{C}$$

satisfying, for all $u, v, w \in V$ and $\alpha, \beta \in \mathbb{C}$:

- **Conjugate symmetry:** $\langle u, v \rangle = \overline{\langle v, u \rangle}$
- **Linearity in the second argument (physics convention):** $\langle u, \alpha v + \beta w \rangle = \alpha \langle u, v \rangle + \beta \langle u, w \rangle$
- **Positive definiteness:** $\langle v, v \rangle \geq 0$, with equality if and only if $v = 0$

A pair $(V, \langle \cdot, \cdot \rangle)$ is called a **complex inner product space**.

The **induced norm** is

$$\|v\| := \sqrt{\langle v, v \rangle}.$$

Two vectors u, v are **orthogonal**, written $u \perp v$, if $\langle u, v \rangle = 0$. A vector is **normalized** (or a **unit vector**) if $\|v\| = 1$.

Note: physics uses the right-linear convention (linear in the second slot, antilinear in the first). Mathematics texts often use the opposite. Throughout this course we use the physics convention exclusively.

Intuition

An inner product is the complex generalization of the dot product. It measures "how much one vector lies along another," but with two twists:

- The output is complex, not real. The complex phase of $\langle \phi, \psi \rangle$ encodes interference — the distinguishing feature of quantum mechanics.
- Only the squared modulus $|\langle \phi, \psi \rangle|^2$ is directly physical, because only it is basis-independent and real-valued. This will become the Born rule.

The norm $\|v\|$ is the length of v; the inner product $\langle \phi, \psi \rangle$ is the amplitude with which ϕ and ψ overlap.

Structural Role

Inner products are what convert the formal linear structure of a complex vector space into the probabilistic structure of quantum mechanics.

- Every quantum probability is $|\langle \phi | \psi \rangle|^2$.
- Every expectation value is $\langle \psi | A | \psi \rangle$.

- The adjoint operation (Node 1.5) is defined with respect to the inner product.
- The spectral theorem (Node 1.6) requires orthogonality of eigenspaces, which requires an inner product.

Without inner products, the course halts at linear algebra and never reaches physics.

Mathematical Development

Cauchy–Schwarz inequality

For all $u, v \in V$,

$$|\langle u, v\rangle| \le \|u\|\,\|v\|,$$

with equality if and only if u and v are linearly dependent.

Triangle inequality

For all $u, v \in V$,

$$\|u + v\| \le \|u\| + \|v\|.$$

This is a direct consequence of Cauchy–Schwarz.

Parallelogram law

$$\|u + v\|^2 + \|u - v\|^2 = 2\|u\|^2 + 2\|v\|^2.$$

This identity characterizes norms that come from an inner product: a norm on a complex vector space is induced by an inner product if and only if it satisfies the parallelogram law.

Polarization identity

The inner product can be recovered from the norm:

$$\langle u, v\rangle = \tfrac{1}{4}\sum_{k=0}^{3} i^k \|u + i^k v\|^2.$$

Inner product and norm carry the same information — the norm just packages it with the phase discarded.

Worked Example

Qubit inner products

Work on $V = \mathbb{C}^2$ with the standard inner product $\langle u, v\rangle = \sum_i \overline{u_i} v_i$.

Computational basis:

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

Directly:

$$\langle 0|0\rangle = 1, \quad \langle 1|1\rangle = 1, \quad \langle 0|1\rangle = 0.$$

Hadamard basis:

$$|+\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \qquad |-\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

Compute:

$$\langle +|-\rangle = \tfrac{1}{2}(\langle 0|0\rangle - \langle 0|1\rangle + \langle 1|0\rangle - \langle 1|1\rangle) = \tfrac{1}{2}(1 - 0 + 0 - 1) = 0.$$

So $\{|+\rangle, |-\rangle\}$ is orthonormal. Now a cross-basis overlap:

$$\langle 0|+\rangle = \tfrac{1}{\sqrt{2}}(\langle 0|0\rangle + \langle 0|1\rangle) = \tfrac{1}{\sqrt{2}}.$$

Its squared modulus is $|\langle 0|+\rangle|^2 = \frac{1}{2}$ — the Born-rule probability of obtaining outcome 0 when measuring a qubit prepared in state $|+\rangle$ in the computational basis.

Cross-Links

- **Linear Algebra:** Inner products and orthogonality — same axiom system over $\mathbb{R}$ or $\mathbb{C}$.
- **Statistics:** Inner products on L^2 underlie covariance, correlation, and conditional expectation.
- **VQA:** Cost functions have the form $\langle\psi(\theta)|H|\psi(\theta)\rangle$ — a single inner product.
- **Quantum Information:** Fidelity $F(\psi, \phi) = |\langle\psi|\phi\rangle|^2$ is an inner product quantity.

Structural Compression

Invariant: A conjugate-symmetric, right-linear, positive-definite pairing $V \times V \to \mathbb{C}$.

Mechanism: Three axioms (conjugate symmetry, linearity, positivity) induce a norm, a geometry, and a probability.

Cross-System Echo: Real inner products give classical Euclidean geometry. The complex generalization retains lengths and angles but adds phase — the structural ingredient that produces interference.

Exercises

1. Verify the three inner-product axioms for the standard inner product $\langle u, v\rangle = \sum_i \overline{u_i} v_i$ on $\mathbb{C}^n$.
2. Prove Cauchy–Schwarz by considering $\|u - \alpha v\|^2 \geq 0$ for an optimally chosen α.
3. Verify the parallelogram law on $\mathbb{C}^2$ for $u = |0\rangle, v = |+\rangle$.
4. Show that the polarization identity recovers $\langle u, v\rangle$ from the norm on $\mathbb{C}^2$, using explicit vectors.
5. Let $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. Show that $\|\psi\| = 1 \iff |\alpha|^2 + |\beta|^2 = 1$, and interpret the result as a normalization condition on a quantum state.
6. Explain, structurally, why only $|\langle\phi|\psi\rangle|^2$ (not $\langle\phi|\psi\rangle$ itself) appears in measurable quantities.

Hilbert Space — Completeness

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Hilbert Spaces **Node:** 1.3

This node is the structural promotion of "complex inner product space" to "Hilbert space." In finite dimensions the distinction is vacuous; in infinite dimensions it is what makes quantum mechanics analytically well-posed.

Formal Definition

A **Hilbert space** $\mathcal{H}$ is a complex inner product space that is **complete** in the metric induced by its norm: every Cauchy sequence in $\mathcal{H}$ converges to a limit in $\mathcal{H}$.

Explicitly: a sequence $\{v_n\} \subset \mathcal{H}$ is **Cauchy** if for every $\varepsilon > 0$ there exists N such that

$$\|v_n - v_m\| < \varepsilon \quad \text{for all } n, m \geq N.$$

$\mathcal{H}$ is complete if every such sequence has a limit $v \in \mathcal{H}$ with $\|v_n - v\| \to 0$.

A Hilbert space is **separable** if it contains a countable dense subset. All Hilbert spaces arising in ordinary quantum mechanics are separable.

Intuition

Completeness is the assertion that "limits live in the space." A rational-number line has holes ($\sqrt{2}$ is missing); the real line does not. In the same way, an inner product space that is not complete has sequences whose "intended limit" is outside the space — physically unacceptable, because nothing prevents you from constructing such a sequence as a limit of experimentally realizable states.

The upgrade from "complex inner product space" to "Hilbert space" is the analytic closure that lets you take limits of superpositions, integrals of wavefunctions, and spectral decompositions of unbounded operators without leaving the state space.

In finite dimensions, every inner product space is automatically complete, so the distinction is invisible. In infinite dimensions — the setting for position, momentum, and field operators — it is essential.

Structural Role

Hilbert space is the arena in which all of quantum mechanics takes place. The state postulate (Node 2.1) identifies physical states with unit rays in a Hilbert space; the evolution postulate requires unitary operators, whose definition presumes an inner product and completeness; the measurement postulate demands spectral decomposition, which requires completeness to generalize beyond finite dimensions.

Every downstream object — density matrices (5.1), quantum channels (7.3), partial traces (5.2), Lindblad generators (8.2) — is built on Hilbert space.

Mathematical Development

The three canonical examples

Finite-dimensional: $\mathbb{C}^n$ with $\langle u, v\rangle = \sum_i \overline{u_i} v_i$. Automatically complete.

Countably infinite-dimensional: $\ell^2(\mathbb{N})$, the space of square-summable complex sequences:

$$\ell^2 = \left\{(a_1, a_2, \ldots) \in \mathbb{C}^{\mathbb{N}} : \sum_{n=1}^{\infty} |a_n|^2 < \infty\right\}$$

with $\langle a, b\rangle = \sum_n \overline{a_n} b_n$.

Continuous: $L^2(\mathbb{R})$, the space of square-integrable complex-valued functions modulo almost-everywhere equality:

$$L^2(\mathbb{R}) = \left\{\psi : \mathbb{R} \to \mathbb{C} : \int_{-\infty}^{\infty} |\psi(x)|^2 \, dx < \infty\right\}$$

with $\langle \phi, \psi\rangle = \int_{-\infty}^{\infty} \overline{\phi(x)} \psi(x) \, dx$.

The last two are separable, infinite-dimensional, and complete — the non-trivial cases.

Separability and basis existence

Every separable Hilbert space admits a countable orthonormal basis (Node 1.4). In non-separable Hilbert spaces an orthonormal basis still exists (by Zorn), but it need not be countable. Separability is what keeps quantum mechanics tractable.

Finite vs infinite dimensions

Property	Finite-dim	Infinite-dim
Completeness automatic	yes	no — must be imposed
Every linear operator bounded	yes	no
Spectral theorem in elementary form	yes	replaced by spectral measure
Position/momentum operators	n/a	unbounded, densely defined

The finite-dimensional case is sufficient for qubit and finite-level physics. The infinite-dimensional case is necessary as soon as position, momentum, or field degrees of freedom appear.

Worked Example

The qubit Hilbert space

For a single qubit, $\mathcal{H} = \mathbb{C}^2$. Two complex dimensions. Every sequence of unit vectors $|\psi_n\rangle$ with $\||\psi_n\rangle - |\psi_m\rangle\| \to 0$ has a limit in $\mathbb{C}^2$ — there is no analytic subtlety.

The particle-on-a-line Hilbert space

For a non-relativistic particle on $\mathbb{R}$, $\mathcal{H} = L^2(\mathbb{R})$. Elements are wavefunctions $\psi(x)$ with $\int |\psi|^2 \, dx < \infty$. The Gaussian $\psi_\sigma(x) = (2\pi\sigma^2)^{-1/4} e^{-x^2/(4\sigma^2)}$ is in $\mathcal{H}$ for every $\sigma > 0$. A sequence of Gaussians with $\sigma_n \to \sigma_\infty > 0$ has a Gaussian limit, which is again in $\mathcal{H}$. A sequence with $\sigma_n \to 0$ has no limit in L^2 — the pointwise limit is a Dirac delta, which is not square-integrable. The latter is why $\delta(x - x_0)$ is not a physical state; it lives outside the Hilbert space, and position "eigenstates" are a formal device (rigged Hilbert space), not genuine vectors.

Cross-Links

- **Linear Algebra:** Vector spaces with inner products — the finite-dimensional case of this node.
- **Statistics:** L^2 is the Hilbert space of square-integrable random variables; conditional expectation is Hilbert-space projection.
- **VQA:** An n-qubit register lives in $(\mathbb{C}^2)^{\otimes n}$ — a finite-dimensional Hilbert space of dimension 2^n.
- **Math-Intuition:** Completeness in analysis — the same idea specialized to the complex-inner-product setting.

Structural Compression

Invariant: Cauchy-completeness in the inner-product norm.

Mechanism: Topological closure of an inner product space under limits of Cauchy sequences.

Cross-System Echo: Banach completeness in classical analysis generalizes to arbitrary norms; Hilbert completeness is the special case with an inner product structure, enabling orthogonality and spectral theory.

Exercises

1. Prove directly that $\mathbb{C}^n$ is complete, using that $\mathbb{C}$ is complete and that convergence in $\mathbb{C}^n$ reduces to coordinate-wise convergence.
2. Give an example of an incomplete inner product space. (Hint: consider the space of eventually-zero sequences as a subspace of ℓ^2.)
3. Show that ℓ^2 is complete. Show that it is separable by exhibiting a countable dense subset.
4. Show that the Dirac delta $\delta(x)$ is not an element of $L^2(\mathbb{R})$.
5. Let $\mathcal{H}$ be separable and $\{e_n\}$ an orthonormal basis. Show that the map $\mathcal{H} \to \ell^2$, $v \mapsto (\langle e_n, v\rangle)$, is an isometric isomorphism. Conclude that there is essentially only one separable infinite-dimensional Hilbert space.
6. Explain structurally why position eigenstates are "formal" rather than genuine elements of the Hilbert space, and why this does not prevent their use in calculations.

Orthonormal Bases and Dirac Notation

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Hilbert Spaces **Node:** 1.4

This node introduces the bra-ket notation that the rest of the course uses without further comment. It also establishes the resolution of identity, the structural identity that makes basis expansions tractable and is the engine of every spectral decomposition and measurement computation.

Formal Definition

Let $\mathcal{H}$ be a separable Hilbert space.

An **orthonormal set** is a collection $\{|e_i\rangle\} \subset \mathcal{H}$ satisfying

$$\langle e_i|e_j\rangle = \delta_{ij}.$$

An orthonormal set is an **orthonormal basis** if it is additionally **complete**: its linear span is dense in $\mathcal{H}$, equivalently,

$$\sum_i |e_i\rangle\langle e_i| = I,$$

where the sum converges in the strong operator topology (in finite dimensions, as an ordinary sum).

Dirac notation

Vectors in $\mathcal{H}$ are written as **kets**: $|\psi\rangle \in \mathcal{H}$.

The **dual space** $\mathcal{H}^*$ consists of continuous linear functionals on $\mathcal{H}$. By the **Riesz representation theorem**, every $f \in \mathcal{H}^*$ has a unique representer $\phi \in \mathcal{H}$ such that $f(\psi) = \langle \phi, \psi\rangle$. The dual element associated with $|\phi\rangle$ is written as a **bra**: $\langle\phi| \in \mathcal{H}^*$.

The pairing is the inner product:

$$\langle\phi|\psi\rangle \equiv \langle\phi, \psi\rangle.$$

A bra next to a ket is a scalar; a ket next to a bra is an operator:

$$|\phi\rangle\langle\psi| : |\chi\rangle \mapsto |\phi\rangle\, \langle\psi|\chi\rangle.$$

Intuition

Dirac notation is a disciplined typographic device that automatically tracks “row vs column” and “vector vs linear-functional”:

- Ket $|\psi\rangle$: column vector, element of $\mathcal{H}$.
- Bra $\langle\psi|$: row vector, element of $\mathcal{H}^*$.
- Bra-ket $\langle\phi|\psi\rangle$: scalar, the inner product.
- Ket-bra $|\phi\rangle\langle\psi|$: rank-one operator.

The notation enforces the typing: any expression that can be written down in Dirac notation is syntactically valid, and its “type” (scalar, vector, operator) is visible from the bracket pattern.

An orthonormal basis is a coordinate system on $\mathcal{H}$. The completeness relation $\sum_i |e_i\rangle\langle e_i| = I$ is what lets you insert a basis expansion anywhere in a computation, the way you insert “1” in an arithmetic identity.

Structural Role

The resolution of identity

$$I = \sum_i |e_i\rangle\langle e_i|$$

is the engine behind every basis expansion, every measurement probability, and every spectral decomposition. Given any state $|\psi\rangle$,

$$|\psi\rangle = I|\psi\rangle = \sum_i |e_i\rangle\langle e_i|\psi\rangle,$$

so the coefficients in the basis expansion are exactly the inner products $\langle e_i|\psi\rangle$. These coefficients, squared, are Born-rule probabilities (Node 2.3). The resolution of identity therefore structurally contains all of measurement statistics as a special case.

Mathematical Development

Basis expansion

For any $|\psi\rangle \in \mathcal{H}$:

$$|\psi\rangle = \sum_i c_i|e_i\rangle, \qquad c_i = \langle e_i|\psi\rangle.$$

Parseval’s identity

$$\|\psi\|^2 = \sum_i |c_i|^2 = \sum_i |\langle e_i|\psi\rangle|^2.$$

This is the Hilbert-space Pythagorean theorem: the squared length of a vector is the sum of the squared moduli of its basis coefficients.

Inner product in coordinates

$$\langle\phi|\psi\rangle = \sum_i \overline{\langle e_i|\phi\rangle}\,\langle e_i|\psi\rangle = \sum_i \langle\phi|e_i\rangle\langle e_i|\psi\rangle.$$

Change of basis

Let $\{|e_i\rangle\}$ and $\{|f_j\rangle\}$ be two orthonormal bases. The change-of-basis operator is

$$U = \sum_j |f_j\rangle\langle e_j|,$$

which satisfies $U^\dagger U = UU^\dagger = I$ — a unitary operator (Node 1.5). A state's coefficients transform as

$$\langle f_j|\psi\rangle = \sum_i \langle f_j|e_i\rangle\langle e_i|\psi\rangle,$$

i.e., multiplication by the matrix $U_{ji} = \langle f_j|e_i\rangle$.

Gram–Schmidt

Any countable linearly independent set $\{v_n\}$ in $\mathcal{H}$ can be converted into an orthonormal set by the Gram–Schmidt procedure:

$$|u_n\rangle = |v_n\rangle - \sum_{k<n} |e_k\rangle\langle e_k|v_n\rangle, \qquad |e_n\rangle = \frac{|u_n\rangle}{\|u_n\|}.$$

This is the constructive proof that every separable Hilbert space has an orthonormal basis.

Worked Example

Computational and Hadamard bases on $\mathbb{C}^2$

Work on the qubit Hilbert space $\mathbb{C}^2$. The **computational basis** is

$$|0\rangle = \begin{pmatrix}1\\0\end{pmatrix}, \qquad |1\rangle = \begin{pmatrix}0\\1\end{pmatrix}.$$

Verify the completeness relation:

$$|0\rangle\langle 0| + |1\rangle\langle 1| = \begin{pmatrix}1 & 0\\0 & 0\end{pmatrix} + \begin{pmatrix}0 & 0\\0 & 1\end{pmatrix} = \begin{pmatrix}1 & 0\\0 & 1\end{pmatrix} = I.$$

The **Hadamard basis** is

$$|+\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \qquad |-\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

Change of basis:

$$|0\rangle = \tfrac{1}{\sqrt{2}}(|+\rangle + |-\rangle), \qquad |1\rangle = \tfrac{1}{\sqrt{2}}(|+\rangle - |-\rangle).$$

The change-of-basis operator is the **Hadamard matrix**

$$H = |+\rangle\langle 0| + |-\rangle\langle 1| = \tfrac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\1 & -1\end{pmatrix},$$

which satisfies $H^\dagger H = I$ and $H^2 = I$.

Expanding a state in two bases

Consider $|\psi\rangle = \frac{1}{\sqrt{3}}|0\rangle + \sqrt{\frac{2}{3}}|1\rangle$. In the computational basis the coefficients are $(1/\sqrt{3}, \sqrt{2/3})$, with squared moduli $(1/3, 2/3)$ — the Born probabilities for outcomes 0 and 1 in a computational measurement.

Hadamard coefficients:

$$\langle +|\psi\rangle = \frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{3}} + \sqrt{\frac{2}{3}}\right), \quad \langle -|\psi\rangle = \frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{3}} - \sqrt{\frac{2}{3}}\right).$$

Parseval: $|\langle +|\psi\rangle|^2 + |\langle -|\psi\rangle|^2 = 1$. The probabilities of the two outcomes depend on the measurement basis, but their sum is invariant — this is the structural reason probability is conserved under change of basis.

Cross-Links

- **Linear Algebra:** Orthonormal bases and Gram–Schmidt.
- **Quantum Information:** Computational vs Hadamard bases are the standard measurement bases for protocols (e.g., BB84).
- **VQA:** Measurement expansion in the computational basis — every cost-function estimate is a sum over basis-projector expectation values.
- **Statistics:** Fourier basis on $L^2(S^1)$ as an example of an infinite-dimensional orthonormal basis; characteristic functions.

Structural Compression

Invariant: The resolution of identity $\sum_i |e_i\rangle\langle e_i| = I$.

Mechanism: Orthonormality plus completeness — two-part axiom.

Cross-System Echo: Fourier expansion in L^2; character sums in harmonic analysis; principal component decomposition in statistics. All are instances of "expand in an orthonormal basis that diagonalizes the structure of interest."

Exercises

1. Verify that $\{|+\rangle, |-\rangle\}$ is an orthonormal basis of $\mathbb{C}^2$: check $\langle +|+\rangle = \langle -|-\rangle = 1$, $\langle +|-\rangle = 0$, and the completeness relation $|+\rangle\langle +| + |-\rangle\langle -| = I$.
2. Prove Parseval's identity from the resolution of identity.
3. Apply Gram–Schmidt to the vectors $v_1 = (1,0)^T$ and $v_2 = (1,1)^T$ in $\mathbb{C}^2$. Check that the result is orthonormal.
4. Show that the Hadamard matrix H satisfies $H = H^\dagger$ and $H^2 = I$. Interpret: H is both a change of basis and its own inverse.
5. For $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, compute the coefficients in the Hadamard basis. Verify that the squared moduli sum to $|\alpha|^2 + |\beta|^2$.
6. Explain, structurally, why the completeness relation is the single identity that generates both basis expansions and measurement probabilities.

Linear Operators and Adjoints

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Hilbert Spaces **Node:** 1.5

The objects acting on quantum states are linear operators. The adjoint operation, defined through the inner product, distinguishes the special operator classes — Hermitian (observables), unitary (evolution), and projection (measurement) — that the postulates of quantum mechanics require.

Formal Definition

Let $\mathcal{H}$ be a Hilbert space.

A **linear operator** on $\mathcal{H}$ is a map $A : \mathcal{H} \to \mathcal{H}$ satisfying

$$A(\alpha|\psi\rangle + \beta|\phi\rangle) = \alpha A|\psi\rangle + \beta A|\phi\rangle$$

for all $|\psi\rangle, |\phi\rangle \in \mathcal{H}$ and $\alpha, \beta \in \mathbb{C}$.

The operator A is **bounded** if there exists $C \geq 0$ with $\|A|\psi\rangle\| \leq C\|\psi\|$ for all $|\psi\rangle$. The smallest such C is the **operator norm** $\|A\|$. In finite dimensions every linear operator is bounded; in infinite dimensions, operators like position and momentum are unbounded and only densely defined.

The **adjoint** $A^\dagger$ of a bounded linear operator A is the unique bounded linear operator satisfying

$$\langle\phi|A\psi\rangle = \langle A^\dagger\phi|\psi\rangle \quad \text{for all } |\phi\rangle, |\psi\rangle \in \mathcal{H}.$$

Existence and uniqueness follow from the Riesz representation theorem applied to the functional $|\psi\rangle \mapsto \langle\phi|A\psi\rangle$.

Operator classes

- **Hermitian (self-adjoint):** $A = A^\dagger$. Corresponds to observables.
- **Unitary:** $U^\dagger U = UU^\dagger = I$. Corresponds to physical evolution and change of basis.
- **Projection (orthogonal):** $P^2 = P$ and $P^\dagger = P$. Corresponds to yes/no measurement questions.
- **Normal:** $[A, A^\dagger] = 0$. The largest class admitting the spectral theorem.

Intuition

Operators are "things that do something to states." They are the verbs of quantum mechanics; states are the nouns.

The adjoint is the natural generalization of the matrix conjugate-transpose. In finite dimensions, if A is represented in an orthonormal basis by the matrix $A_{ij} = \langle e_i|A|e_j\rangle$, then $A^\dagger_{ij} = \overline{A_{ji}}$. Conjugate, then transpose.

The three distinguished classes correspond to the three things one does with a quantum state:

- **Hermitian $\leftrightarrow$ real eigenvalues $\leftrightarrow$ measurable quantities.** Eigenvalues are the possible outcomes of a measurement, and measurement outcomes are real numbers.
- **Unitary $\leftrightarrow$ norm-preserving $\leftrightarrow$ probability-conserving evolution.** Unitaries are invertible and preserve inner products, which is what makes Born-rule probabilities a conserved total.
- **Projection $\leftrightarrow$ idempotent yes/no $\leftrightarrow$ measurement outcomes as subspaces.** A projection partitions the Hilbert space into "inside the subspace" and "orthogonal to it" — the two possibilities of a yes/no test.

Structural Role

The three operator classes each carry a postulate of quantum mechanics:

- Hermitian operators represent observables (Postulate 2, Node 2.2).
- Unitary operators represent closed-system time evolution (Postulate 4, Node 2.4).
- Projections (and their generalizations, POVMs) represent measurements (Postulates 3, 3.1, 3.2).

The adjoint is the single algebraic device that organizes these distinctions: all three classes are defined by adjoint equations. Without the inner product (Node 1.2), none of these definitions exist.

Downstream, the adjoint is also the structural origin of quantum channels (Node 7.3), because Kraus operators arise as pairs $\{K_i\}$ with $\sum_i K_i^\dagger K_i = I$ — an adjoint condition that guarantees probability conservation.

Mathematical Development

Adjoint algebra

For bounded operators A, B and scalar α:

$$(A+B)^\dagger = A^\dagger + B^\dagger, \quad (\alpha A)^\dagger = \overline{\alpha} A^\dagger, \quad (AB)^\dagger = B^\dagger A^\dagger, \quad (A^\dagger)^\dagger = A.$$

Note the order reversal in $(AB)^\dagger$ and the complex conjugation of scalars.

Trace

For a trace-class operator, the **trace** is

$$\mathrm{Tr}(A) = \sum_i \langle e_i | A | e_i \rangle,$$

independent of the orthonormal basis $\{|e_i\rangle\}$. Key properties:

$$\mathrm{Tr}(AB) = \mathrm{Tr}(BA), \quad \mathrm{Tr}(A^\dagger) = \overline{\mathrm{Tr}(A)}.$$

The trace will reappear in Node 5.1 as the normalization condition for density matrices.

Commutators

The **commutator** $[A, B] := AB - BA$ measures non-commutativity. For Hermitian A, B:

$$[A, B]^\dagger = -[A, B],$$

so commutators of Hermitians are anti-Hermitian, and $i[A, B]$ is Hermitian. Commutators generate the uncertainty relations of Node 3.5.

Characterizations of the three classes

- **Hermitian** $\iff$ $\langle \psi | A | \psi \rangle \in \mathbb{R}$ for all $|\psi\rangle$. The expectation value of a Hermitian operator is real, which is why observables must be Hermitian.
- **Unitary** $\iff$ $\| U|\psi\rangle \| = \|\psi\|$ for all $|\psi\rangle$. Norm preservation is equivalent to the adjoint condition $U^\dagger U = I$.
- **Projection** $\iff$ P is the orthogonal projection onto some closed subspace $\mathcal{M} \subseteq \mathcal{H}$.

Rank-one operators

The ket-bra $|\phi\rangle\langle\psi|$ is a rank-one operator with

$$(|\phi\rangle\langle\psi|)^\dagger = |\psi\rangle\langle\phi|.$$

In particular, $|\psi\rangle\langle\psi|$ is a Hermitian rank-one projection onto the one-dimensional subspace spanned by $|\psi\rangle$ (provided $\|\psi\| = 1$).

Worked Example

The Pauli matrices on $\mathbb{C}^2$

The Pauli operators are

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Hermiticity. Compute $X^\dagger$ by conjugate-transposing: $X^\dagger = X$. Similarly $Y^\dagger = Y$ (note: the off-diagonal entries are $-i$ and i, which swap and conjugate back to themselves) and $Z^\dagger = Z$. All three are Hermitian.

Unitarity. Each Pauli matrix satisfies $X^2 = Y^2 = Z^2 = I$, so $X^\dagger X = X^2 = I$, and similarly for Y and Z. All three are unitary. (A Hermitian operator that squares to the identity is automatically unitary — an uncommon but structurally important coincidence.)

Algebra. Direct computation:

$$[X, Y] = XY - YX = 2iZ, \quad [Y, Z] = 2iX, \quad [Z, X] = 2iY.$$

These are the defining relations of $\mathfrak{su}(2)$, the Lie algebra of rotations — the algebraic source of spin-$\frac{1}{2}$ behavior.

Spectral data. Z has eigenvalues ± 1 with eigenvectors $|0\rangle, |1\rangle$. X has eigenvalues ± 1 with eigenvectors $|\pm\rangle$. Y has eigenvalues ± 1 with eigenvectors $|\pm i\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm i|1\rangle)$.

A projection example

The operator $P_+ = |0\rangle\langle 0|$ satisfies

$$P_+^2 = |0\rangle\langle 0|0\rangle\langle 0| = |0\rangle\langle 0| = P_+, \quad P_+^\dagger = (|0\rangle\langle 0|)^\dagger = |0\rangle\langle 0| = P_+.$$

It projects onto the one-dimensional subspace $\mathrm{span}\{|0\rangle\}$. Its complement $P_- = I - P_+ = |1\rangle\langle 1|$ is also a projection, and $P_+ + P_- = I$. Together they form the computational-basis measurement (Node 3.1).

Cross-Links

- **Linear Algebra:** Linear transformations and adjoints (module on inner product spaces). Matrix transpose is the real special case.
- **Statistics:** Self-adjoint operators on L^2 generalize the notion of random variable; covariance operators are Hermitian positive-semidefinite.
- **VQA:** Hamiltonians are Hermitian observables; variational cost functions are inner products $\langle\psi(\theta)|H|\psi(\theta)\rangle$.

- **Quantum Information:** Kraus operators of quantum channels are constrained by $\sum_i K_i^\dagger K_i = I$ — an adjoint relation.

Structural Compression

Invariant: Linearity, plus the adjoint equation $\langle\phi|A\psi\rangle = \langle A^\dagger\phi|\psi\rangle$.

Mechanism: One linearity axiom and one duality relation (via Riesz) carve out the three operator classes needed by the quantum postulates.

Cross-System Echo: Matrix transpose is the real special case of the adjoint; orthogonal matrices are the real analogue of unitaries; symmetric matrices are the real analogue of Hermitians. Quantum mechanics is the complex-linear upgrade of this classical structure.

Exercises

1. Show that the adjoint is unique: if B_1 and B_2 both satisfy $\langle\phi|A\psi\rangle = \langle B_k\phi|\psi\rangle$ for all $|\phi\rangle, |\psi\rangle$, then $B_1 = B_2$.
2. Prove that $(AB)^\dagger = B^\dagger A^\dagger$ directly from the defining relation of the adjoint.
3. Prove that the expectation value of a Hermitian operator is real: $\langle\psi|A|\psi\rangle \in \mathbb{R}$ when $A = A^\dagger$.
4. Prove that a linear operator is unitary if and only if it preserves the norm of every vector.
5. Verify the Pauli commutation relations $[X, Y] = 2iZ$, $[Y, Z] = 2iX$, $[Z, X] = 2iY$ by direct matrix calculation.
6. Show that the ket-bra $|\psi\rangle\langle\psi|$ (with $\|\psi\| = 1$) is a projection. Compute its trace and its action on an arbitrary vector.
7. Let A be Hermitian and U unitary. Show that $UAU^\dagger$ is Hermitian and has the same eigenvalues as A. Interpret this as “change of basis preserves observables.”

Spectral Decomposition

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Hilbert Spaces **Node:** 1.6

The spectral theorem is the bridge from linear algebra into quantum measurement. Every observable is decomposed into a sum over eigenvalues weighted by orthogonal projections — and this decomposition is exactly what the Born rule (Node 2.3) and projective measurement (Node 3.1) consume. This is the closing node of Module 1; the next module will start converting this machinery into physics.

Formal Definition

Finite-dimensional spectral theorem

Let $\mathcal{H}$ be a finite-dimensional complex Hilbert space and A a Hermitian operator on $\mathcal{H}$. Then A admits a **spectral decomposition**

$$A = \sum_i \lambda_i P_i,$$

where:

- $\lambda_i \in \mathbb{R}$ are the distinct eigenvalues of A,
- P_i is the orthogonal projection onto the eigenspace $\mathcal{H}_i = \ker(A - \lambda_i I)$,
- the projections are mutually orthogonal: $P_i P_j = \delta_{ij} P_i$,
- the projections resolve the identity: $\sum_i P_i = I$.

The eigenspaces give an orthogonal decomposition $\mathcal{H} = \bigoplus_i \mathcal{H}_i$. The theorem holds more generally for **normal** operators (those satisfying $[A, A^\dagger] = 0$), with complex eigenvalues; Hermiticity is the special case that guarantees real eigenvalues.

Spectral theorem (general form)

For a bounded normal operator A on a (possibly infinite-dimensional) Hilbert space, the spectral sum is replaced by a **spectral measure**: a projection-valued measure E on the complex plane (supported on the spectrum of A) such that

$$A = \int_{\sigma(A)} \lambda \, dE(\lambda).$$

The discrete-sum version is recovered when the spectrum is a countable set of isolated eigenvalues.

Intuition

Every observable is a list of real outcomes plus a partition of Hilbert space into "what would collapse to this outcome" subspaces.

Think of the Hermitian operator A as a compound object with two pieces of information:

1. A finite list of real numbers $\lambda_1, \lambda_2, \ldots$ — the possible measurement outcomes.
2. A matching list of orthogonal subspaces — one per outcome — that jointly tile the whole Hilbert space.

A state "lying inside" eigenspace i deterministically yields outcome λ_i. A state that is a superposition across eigenspaces yields outcome i with probability equal to the squared length of its projection onto that eigenspace. That is the Born rule, prefigured by the spectral theorem.

The spectral theorem is why quantum mechanics is diagonal-in-the-right-basis: for any observable, there is a basis (its eigenbasis) in which measurement is trivial to compute.

Structural Role

The spectral theorem is the structural hinge of the course. Everything before it is abstract linear algebra; everything after it is physics.

- It supplies the projections that appear in projective measurement (Node 3.1).
- It guarantees that every observable has real outcomes (Node 2.2).
- It gives the Born rule an algebraic shape: $\Pr(\lambda_i) = \mathrm{Tr}(P_i\rho)$ (Node 2.3, Node 5.1).
- It diagonalizes density matrices into mixtures of pure states (Node 5.1).
- It diagonalizes Hamiltonians into time-evolution-generating eigenmodes (Node 2.4).
- In infinite dimensions, it is what gives position, momentum, and energy their spectra and their associated projective measures.

No spectral theorem, no measurement theory.

Mathematical Development

Proof sketch (finite-dimensional)

Let A be Hermitian on $\mathcal{H}$ with $\dim\mathcal{H} = n < \infty$.

Step 1 — existence of a real eigenvalue. The characteristic polynomial $\det(A - \lambda I)$ has at least one root in $\mathbb{C}$. For any eigenvector $|v\rangle$ with $A|v\rangle = \lambda|v\rangle$ and $\|v\| = 1$:

$$\lambda = \langle v|A|v\rangle = \overline{\langle v|A|v\rangle} = \overline{\lambda},$$

using $A = A^\dagger$. So every eigenvalue is real.

Step 2 — eigenspaces are orthogonal. If $A|v\rangle = \lambda|v\rangle$ and $A|w\rangle = \mu|w\rangle$ with $\lambda \neq \mu$, then

$$\lambda\langle w|v\rangle = \langle w|Av\rangle = \langle Aw|v\rangle = \mu\langle w|v\rangle,$$

forcing $\langle w|v\rangle = 0$.

Step 3 — induction on dimension. Let $\mathcal{H}_\lambda = \ker(A - \lambda I)$ be the eigenspace for one eigenvalue λ. Its orthogonal complement $\mathcal{H}_\lambda^\perp$ is invariant under A: for $|w\rangle \in \mathcal{H}_\lambda^\perp$ and $|v\rangle \in \mathcal{H}_\lambda$,

$$\langle v|Aw\rangle = \langle Av|w\rangle = \lambda\langle v|w\rangle = 0,$$

so $A|w\rangle \in \mathcal{H}_\lambda^\perp$. Restrict A to $\mathcal{H}_\lambda^\perp$, which is strictly lower-dimensional, and apply the induction hypothesis.

Step 4 — assemble. Iterating gives an orthogonal decomposition $\mathcal{H} = \bigoplus_i \mathcal{H}_i$. Set P_i to be the orthogonal projection onto $\mathcal{H}_i$. Then $P_iP_j = \delta_{ij}P_i$, $\sum_i P_i = I$, and $A = \sum_i \lambda_i P_i$ because both sides agree on a basis.

Functional calculus

Once $A = \sum_i \lambda_i P_i$, one defines functions of operators by applying the function to eigenvalues:

$$f(A) := \sum_i f(\lambda_i) P_i.$$

This gives exponentials $e^{-iHt/\hbar} = \sum_i e^{-i\lambda_i t/\hbar} P_i$ (the time-evolution operator of Node 2.4), powers, square roots, and logarithms — all constructed directly from the spectral data.

Unitary diagonalization

Equivalently: every Hermitian A can be written as $A = UDU^\dagger$ where D is diagonal with the eigenvalues of A on the diagonal and U is the unitary whose columns are the eigenvectors of A. The two forms are equivalent: $P_i = U\Pi_i U^\dagger$ where Π_i is the diagonal projection onto coordinate i in the standard basis.

Worked Example

Pauli Z

$$Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Eigenvalues: $+1$ with eigenvector $|0\rangle$, and -1 with eigenvector $|1\rangle$. Projections: $P_+ = |0\rangle\langle 0|$, $P_- = |1\rangle\langle 1|$. Spectral decomposition:

$$Z = (+1)|0\rangle\langle 0| + (-1)|1\rangle\langle 1|.$$

Verify the resolution of identity: $P_+ + P_- = |0\rangle\langle 0| + |1\rangle\langle 1| = I$. Verify orthogonality: $P_+ P_- = |0\rangle\langle 0|1\rangle\langle 1| = 0$.

Pauli X

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Eigenvalues: $+1$ with eigenvector $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and -1 with eigenvector $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Spectral decomposition:

$$X = (+1)|+\rangle\langle +| + (-1)|-\rangle\langle -|.$$

In matrix form: $|+\rangle\langle +| = \frac{1}{2}\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$, $|-\rangle\langle -| = \frac{1}{2}\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$. Their difference reproduces X; their sum is I.

Functional calculus on Z

Compute e^{-iZt}:

$$e^{-iZt} = e^{-it}|0\rangle\langle 0| + e^{+it}|1\rangle\langle 1| = \begin{pmatrix} e^{-it} & 0 \\ 0 & e^{+it} \end{pmatrix}.$$

This is the time-evolution operator for a qubit Hamiltonian $H = Z$ (taking $\hbar = 1$). The exponential is trivial in the eigenbasis — the spectral theorem is what lets you compute it.

Cross-Links

- **Linear Algebra:** The spectral theorem is the signature theorem of the linear algebra course. This node is its complex-Hermitian specialization.
- **Statistics:** Principal component analysis is the spectral decomposition of a covariance matrix. The same machinery; different interpretation.
- **VQA:** Variational eigensolvers target the smallest eigenvalue of a spectral decomposition via parameterized circuits; the ground state is the $\lambda_{\min}$ eigenvector.
- **Quantum Information:** Von Neumann entropy (Node 7.4) is computed by spectrally decomposing a density matrix.
- **Quantum Thermodynamics:** The Gibbs state $e^{-\beta H}/Z$ is defined by the functional calculus applied to the Hamiltonian.

Structural Compression

Invariant: Every Hermitian operator equals an eigenvalue-weighted sum of mutually orthogonal projections that resolve the identity.

Mechanism: Orthogonality of eigenspaces (from the adjoint relation) plus induction on dimension (finite case), or projection-valued measures (infinite case).

Cross-System Echo: Diagonalization of symmetric matrices in real linear algebra; Fourier analysis as the spectral decomposition of the translation operator; principal components as the spectral decomposition of a covariance operator. All are instances of "resolve an operator into outcomes plus a basis adapted to those outcomes."

Exercises

1. Compute the spectral decomposition of Pauli Y. Write out the projections explicitly as 2×2 matrices.
2. Let A be Hermitian with spectral decomposition $A = \sum_i \lambda_i P_i$. Prove that $A^2 = \sum_i \lambda_i^2 P_i$ and, more generally, $p(A) = \sum_i p(\lambda_i) P_i$ for any polynomial p.
3. Show that $\mathrm{Tr}(A) = \sum_i \lambda_i \dim \mathcal{H}_i$ where $\mathcal{H}_i$ is the eigenspace for λ_i.
4. Let A be Hermitian on $\mathbb{C}^n$ with spectral decomposition $A = \sum_i \lambda_i P_i$. Show that $\langle\psi|A|\psi\rangle = \sum_i \lambda_i \|P_i|\psi\rangle\|^2$ and interpret this as a weighted average of eigenvalues with weights given by Born-rule probabilities.
5. Construct an operator on $\mathbb{C}^3$ with eigenvalues $(0, 1, 1)$ and a non-trivial two-dimensional eigenspace for the eigenvalue 1. Exhibit its spectral decomposition with two (not three) projections.
6. Show that a projection P is itself a Hermitian operator whose spectral decomposition has only the eigenvalues 0 and 1. What are the two eigenspaces?
7. Explain, structurally, why the spectral theorem is the bridge between "linear algebra" and "measurement theory" in this course.