

Geometry as a Realization Functor

Part V of *Toward a Condensed Representation Theory of Physics*

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

07 July 2026

Abstract

A recurring pattern in twentieth- and twenty-first-century mathematics is that geometric objects can be *reconstructed* as invariants of purely algebraic or representation-theoretic data. Tannakian duality recovers an affine group scheme from its tensor category of representations equipped with a fiber functor; Gelfand duality recovers a compact Hausdorff space from its commutative C^* -algebra; Connes’ spectral reconstruction theorem recovers a closed Riemannian spin manifold from a commutative spectral triple; Tomita–Takesaki modular theory together with the Bisognano–Wichmann theorem recovers Lorentz boosts and the causal complement from a von Neumann algebra and a cyclic–separating vector; and the cobordism hypothesis presents a topological field theory as a functor rigidly determined by its value on a point. This paper organizes these five *established* reconstruction theorems into a single meta-pattern — *geometry is the realization of algebraic data* — and casts the underlying arrow as a *realization functor* along a precise representation ladder

abstract algebra $\rightarrow$ condensed object $\rightarrow$ realization $\rightarrow$ observable $\rightarrow$ geometry.

Established mathematics [EST]: the fiber-functor formalism, the Tannakian reconstruction theorem (Saavedra Rivano; Deligne–Milne), Gelfand naturality, Connes’ reconstruction theorem and spectral distance formula, and Bisognano–Wichmann; together with two fully worked finite toy models — Tannakian reconstruction of a finite group from its representation category, and Connes’ spectral distance on a finite spectral triple realized as a graph metric — for which we also supply verified code. *Heuristic dictionary* [HEU]: that a metric can be read off from realization data by a spectral-distance construction, and that the four independent precedents are instances of one structural mechanism. *Speculative hypothesis* [SPEC], clearly boxed: that *physical spacetime itself is the realization of condensed algebraic data* — that a suitably enriched *condensed Tannakian* category of observables, together with a condensed fiber functor, reconstructs an emergent (ultimately Lorentzian, dynamical) geometry. We

mark the boundary precisely, and we are candid about the single sharpest gap: every established reconstruction of a metric from spectral data is *Riemannian*, whereas physical spacetime is *Lorentzian*, and no reconstruction theorem recovering full causal structure from representation data is known.

Contents

1	Introduction	3
1.1	The modular program and its governing perspective	3
1.2	Reconstruction as a pattern: five precedents	3
1.3	The representation ladder	4
1.4	The epistemic discipline: EST, HEU, SPEC	5
1.5	What this Part is and is not	5
1.6	Contributions and outline	6
2	Fiber functors and Tannakian reconstruction	6
2.1	Tensor categories and rigidity	6
2.2	The reconstruction theorem	7
2.3	A worked established example: mixed Tate motives	8
2.4	The finite case, in full	8
3	Four more precedents, one schema	9
3.1	Gelfand duality: a space from a commutative C^* -algebra	9
3.2	Connes' spectral reconstruction: a Riemannian manifold from a spectral triple	10
3.3	Tomita–Takesaki and Bisognano–Wichmann: boosts from a von Neumann algebra	10
3.4	The cobordism hypothesis: physics as a rigidly determined functor	11
3.5	The common schema	11
4	The realization functor and the representation ladder	12
4.1	Realization functors, abstractly	12
4.2	The ladder as a commuting diagram	13
5	Emergent metrics: spectral distance as the EST toy model	14
5.1	Finite spectral triples and the spectral distance	14
5.2	The two-point space: a metric from a single off-diagonal entry	15
5.3	Graph metrics as spectral distances	15
6	The speculative core: condensed reconstruction of spacetime	16
6.1	The hypothesis	17
6.2	What is established, what is proposed	17
6.3	The Riemannian–Lorentzian gap, honestly	18
6.4	A concrete research target	19
7	Results	19

8	Discussion	20
8.1	What has and has not been shown	20
8.2	EST/HEU/SPEC census	20
8.3	Relations and non-identifications	21
8.4	Limitations	21
8.5	Outlook: hand-off to Parts VI–VII	21
9	Conclusion	22
10	The accompanying code	22

1 Introduction

1.1 The modular program and its governing perspective

This is Part V of a modular, hierarchical seven-part program organized by a single perspective:

Physics is the study of realizations of condensed mathematical structures.

The program is *modular*, not unified: each Part isolates one interface between condensed mathematics and physics, states precisely what it establishes, and composes with its predecessors only through explicitly declared hand-offs. Part I [1] replaced the smooth manifold by the condensed set as the kinematic substrate, exhibiting a fully faithful embedding $\mathrm{Top}_{\mathrm{cg}} \hookrightarrow \mathrm{Cond}(\mathrm{Set})$ of compactly generated spaces into condensed sets. Part II [2] attached observables to condensed test objects, obtaining a (co)sheaf of observable algebras on the pro-étale site before any geometry is posited. Part III [3] proposed that entanglement is a compatibility (cocycle) condition in condensed cohomology. Part IV [4] demoted the manifold on which a field lives, replacing $\varphi: M \rightarrow \mathbb{R}$ by a condensed-module-valued object $\varphi: C \rightarrow \mathrm{Cond}(\mathrm{Vect})$, and audited exactly which ingredients of the free field still secretly invoke M .

Each of Parts I–IV *removes* a piece of presupposed geometric background. The present Part asks the reciprocal, constructive question: once the manifold has been demoted, *how does geometry come back?* Our answer, at the level of established mathematics, is that geometry returns as the value of a *realization functor* applied to algebraic data — and that this is not a new idea but the common structural core of several of the deepest reconstruction theorems in modern mathematics.

1.2 Reconstruction as a pattern: five precedents

The central observation of this Part is that the following five results, superficially from unrelated fields, all instantiate the same schema:

a geometric or dynamical object is recovered as an invariant of algebraic/representation data, with no such object presupposed in the input.

1. **Gelfand duality** [EST]. A commutative unital C^* -algebra A is isometrically $*$ -isomorphic to $C(X)$ for a compact Hausdorff space $X = \text{Spec}(A)$ (the Gelfand spectrum of characters); the space is *recovered* from the algebra, functorially and naturally [11].
2. **Tannakian reconstruction** [EST]. A neutral Tannakian category $\mathcal{C}$ over a field k , equipped with a fiber functor $\omega: \mathcal{C} \rightarrow \text{Vect}_k^{\text{fd}}$, satisfies $\mathcal{C} \simeq \text{Rep}_k(G_\omega)$ where $G_\omega = \underline{\text{Aut}}^\otimes(\omega)$ is an affine group scheme; the *group scheme is recovered from its representation category* [8, 7].
3. **Connes' spectral reconstruction** [EST]. A commutative spectral triple $(\mathcal{A}, \mathcal{H}, D)$ satisfying Connes' seven axioms is the canonical spectral triple of a closed Riemannian spin manifold, whose geodesic distance is recovered by the *spectral distance formula* $d_D(p, q) = \sup\{|f(p) - f(q)| : \|[D, f]\| \leq 1\}$; the *manifold is recovered from the algebra and the Dirac operator* [12, 13].
4. **Tomita–Takesaki / Bisognano–Wichmann** [EST]. For the von Neumann algebra of a Rindler wedge with the vacuum as cyclic–separating vector, the modular flow $\sigma_t = \Delta^{it}(\cdot)\Delta^{-it}$ coincides with the wedge-preserving Lorentz boosts and $J\mathcal{M}J = \mathcal{M}'$ realizes Haag duality; *geometric boosts and the causal complement are recovered from operator-algebraic data* [15, 16, 17].
5. **Cobordism hypothesis** [EST]. A fully extended topological field theory $Z: \text{Bord}_n \rightarrow \mathcal{C}$ valued in a symmetric monoidal (∞, n) -category with duals is determined by the fully dualizable object $Z(\text{pt})$; the *functor is rigidly reconstructed from minimal categorical data* [18, 19].

Precedents (1)–(3) recover *metric/topological* data; (4) recovers *Lorentzian causal and thermal* data; (5) recovers a *functor* (physics-as-a-functor) from a single object. Read together they license a working slogan — *geometry is a realization functor* — which we make into a precise organizing diagram in Section 4 and then ask, as this Part's central speculative question, whether a *condensed* enrichment of the pattern can reconstruct physical spacetime.

1.3 The representation ladder

The inherited representation-stack formalism of the collaboration [25] presents physical content as a pipeline

$$M \xrightarrow{\Phi} \Phi(M) \xrightarrow{\text{Real}_\alpha} \text{physical representation} \xrightarrow{\text{Obs}} \text{observable},$$

$$\text{Obs}_\alpha(M) = \text{Obs}(\text{Real}_\alpha(\Phi(M))).$$

Part V specializes the middle arrow Real_α — the *Realization axiom* — to the condensed setting and reads it as a fiber/realization functor. The ladder we study is

$$\underbrace{\mathcal{C}}_{\text{abstract algebra}} \xrightarrow{\iota} \underbrace{\mathcal{C}^{\text{Cond}}}_{\text{condensed object}} \xrightarrow{\omega} \underbrace{\text{Cond}(\text{Vect}_k)}_{\text{realization}} \xrightarrow{\langle \cdot, \cdot \rangle} \underbrace{\text{states/values}}_{\text{observable}} \xrightarrow{d} \underbrace{(X, g)}_{\text{geometry}}. \quad (1)$$

The first three arrows are, in the classical (non-condensed) case, exactly the data of a Tannakian category with a fiber functor; the fourth is evaluation of representations on states (observables); the fifth is a distance/metric read-off in the style of Connes. Our contention is that (a) the entire left-hand portion is *established* for the classical target $\text{Vect}_k^{\text{fd}}$, and (b) the physically interesting content of the program lies in enlarging the target to $\text{Cond}(\text{Vect}_k)$ and asking whether the fifth arrow can produce a *Lorentzian* (X, g) .

1.4 The epistemic discipline: EST, HEU, SPEC

We enforce, verbatim in spirit, the collaboration’s three-valued status calculus. Each substantive claim is tagged:

- **[EST]** — *established*: a theorem with a proof in the cited literature, or a routine consequence thereof reproduced here.
- **[HEU]** — *heuristic*: a physically or structurally motivated dictionary entry that is *not* itself a theorem.
- **[SPEC]** — *speculative*: an ontological or physical hypothesis of this program, not established.

Warrants are ordered **[EST]** < **[HEU]** < **[SPEC]** and composition is monotone, *worst-component-wins*: any chain of reasoning is capped at the status of its weakest link [27]. In particular, the mathematics of Sections 2, 3 and 5 is **[EST]**; the meta-pattern reading of Section 4 is **[HEU]**; and the headline claim “spacetime is the realization of condensed data” of Section 6 is **[SPEC]**, and is boxed accordingly. We are especially careful about one non-identification, restated wherever it is relevant: Connes’ reconstruction is *Riemannian and static*; a Lorentzian, dynamical reconstruction recovering causal structure is *open* [14], and we never silently upgrade the former’s rigor to the latter’s claim.

1.5 What this Part is and is not

This Part *is*: (i) an exposition, in one common language, of five rigorous reconstruction theorems as instances of a realization-functor pattern; (ii) a precise formulation of the representation ladder (1) and of what a “condensed fiber functor” would have to be; (iii) two fully worked, code-backed finite toy models — a Tannakian reconstruction of a finite group and a Connes spectral distance realized as a graph metric — that make the abstract arrows concrete and checkable; and (iv) an honest catalogue of the gap between the established Riemannian reconstructions and the desired Lorentzian one.

This Part *is not*: a claim to have constructed a condensed Tannakian category reconstructing spacetime; a proof that any physical metric emerges from condensed data; or a solution to the Lorentzian reconstruction problem. The central generalization — a *condensed* Tannakian reconstruction yielding emergent geometry — is offered as a labeled **[SPEC]** research proposal, not as a theorem. We flag, once and prominently, that no literature known to us constructs a “condensed Tannakian category” or “condensed fiber functor” in the sense this Part contemplates; this is original, and unproven, research territory.

1.6 Contributions and outline

- Section 2 recalls the fiber-functor formalism and the neutral Tannakian reconstruction theorem, with the mixed Tate motives category as a worked [EST] example, and gives a self-contained finite version (Proposition 2.7).
- Section 3 presents Gelfand, Connes, Bisognano–Wichmann, and the cobordism hypothesis side by side, and distills the common *reconstruction schema* (Principle 3.9).
- Section 4 defines the realization functor and the representation ladder (1) as a precise commuting diagram, and states the Realization axiom in condensed form.
- Section 5 develops the emergent-metric side: Connes’ spectral distance as the [EST] toy model of “metric from realization data,” with a complete finite spectral triple / graph-metric worked example (Theorem 5.6).
- Section 6 states the central [SPEC] proposal — condensed Tannakian reconstruction of emergent geometry — as a boxed hypothesis, and dissects the Riemannian-versus-Lorentzian gap.
- Section 7 collects the results; Section 8 gives the EST/HEU/SPEC census, non-identifications, and limitations; Section 9 concludes. Section 10 documents the accompanying Haskell.

2 Fiber functors and Tannakian reconstruction

We recall the established formalism that is the central technical anchor of this Part. Throughout, k is a field and $\mathbf{Vect}_k^{\text{fd}}$ is the category of finite-dimensional k -vector spaces.

2.1 Tensor categories and rigidity

Definition 2.1 (Tensor category). A k -linear tensor category is a k -linear abelian category $\mathcal{C}$ equipped with a k -bilinear symmetric monoidal structure $(\otimes, \mathbf{1}, a, c, \ell, r)$ for which $\otimes$ is exact in each variable and $\text{End}_{\mathcal{C}}(\mathbf{1}) = k$. An object X is *dualizable* (or *rigid*) if there is $X^\vee$ with unit $\eta: \mathbf{1} \rightarrow X \otimes X^\vee$ and counit $\varepsilon: X^\vee \otimes X \rightarrow \mathbf{1}$ satisfying the triangle (zig-zag) identities. The category $\mathcal{C}$ is *rigid* if every object is dualizable.

Rigidity is the categorical trace of finite-dimensionality: in $\mathbf{Vect}_k^{\text{fd}}$ every space has a dual, and the zig-zag identities encode $\dim V = \text{tr}(\text{id}_V) \in k$. A rigid k -linear tensor category with $\text{End}(\mathbf{1}) = k$ has a well-defined *categorical dimension* $\dim X = \text{tr}(\text{id}_X) \in k$ and *trace* $\text{tr}(f) \in k$ for endomorphisms $f: X \rightarrow X$.

Definition 2.2 (Fiber functor; neutral Tannakian category). A *fiber functor* on a rigid k -linear tensor category $\mathcal{C}$ is an exact, faithful, k -linear symmetric monoidal functor $\omega: \mathcal{C} \rightarrow \mathbf{Vect}_k^{\text{fd}}$. A *neutral Tannakian category* over k is a rigid k -linear abelian tensor category $\mathcal{C}$ with $\text{End}(\mathbf{1}) = k$ that admits a fiber functor.

The fiber functor is the abstract shadow of “choosing coordinates”: it forgets the intrinsic structure of $\mathcal{C}$ and records only underlying vector spaces, exactly, functorially, and compatibly with tensor products. The remarkable content of Tannakian duality is that this forgetful datum loses *nothing*: the whole category, and a group scheme, can be recovered from it.

2.2 The reconstruction theorem

Definition 2.3 (Automorphism group scheme of a fiber functor). Let $(\mathcal{C}, \omega)$ be neutral Tannakian over k . For a commutative k -algebra R , let ω_R denote ω post-composed with $- \otimes_k R$, valued in finitely generated projective R -modules. Define the functor

$$\underline{\mathrm{Aut}}^\otimes(\omega): \mathrm{Alg}_k \rightarrow \mathrm{Gpd}, \quad R \mapsto \mathrm{Aut}^\otimes(\omega_R),$$

the group of monoidal natural automorphisms of ω_R (tensor-compatible, R -linear, natural isomorphisms $\omega_R \Rightarrow \omega_R$).

Theorem 2.4 (Tannakian reconstruction; Saavedra Rivano; Deligne–Milne). [EST] *Let $(\mathcal{C}, \omega)$ be a neutral Tannakian category over k . Then:*

1. *The functor $\underline{\mathrm{Aut}}^\otimes(\omega)$ is representable by an affine group scheme $G_\omega = \mathrm{Spec}(A_\omega)$ over k (the Tannakian dual).*
2. *The functor ω lifts to an equivalence of k -linear tensor categories*

$$\mathcal{C} \xrightarrow{\simeq} \mathrm{Rep}_k(G_\omega),$$

where $\mathrm{Rep}_k(G_\omega)$ is the category of finite-dimensional k -linear representations of G_ω , and ω becomes the forgetful functor $\mathrm{Rep}_k(G_\omega) \rightarrow \mathrm{Vect}_k^{\mathrm{fd}}$.

3. *For a second fiber functor ω' valued in R -modules, the functor $R \mapsto \mathrm{Isom}^\otimes(\omega \otimes R, \omega')$ is representable by a G_ω -torsor over $\mathrm{Spec}(R)$; fiber functors valued in R correspond to G_ω -torsors over R .*

Reference. This is the main theorem of neutral Tannakian duality; see Deligne–Milne [7], Theorems 2.11 and 3.2, building on Saavedra Rivano [8]. We reproduce the finite-group specialization below (Proposition 2.7) with a complete elementary argument, and we implement it in code (Section 10). $\square$

Theorem 2.4 is, verbatim, the statement that an *affine group scheme is recovered from its representation category plus a fiber functor* — geometry (a scheme) from algebra (a tensor category). It is the template on which the entire Part is built.

Remark 2.5 (Coordinate-independence and torsors). Part (3) is the structural heart of the “realization” reading. Different fiber functors are different *realizations* (choices of coordinates) of the same abstract category; the scheme of comparisons between two realizations is a G_ω -torsor. This is precisely the collaboration’s *Equivalence axiom* [25]: internally-definable realizations cannot distinguish equivalent structures, and the ambiguity between realizations is exactly a torsor under the reconstructed symmetry. The motivic incarnation (below) makes this concrete: Betti and de Rham realizations of a motive differ by a period matrix, a point of a torsor under the motivic Galois group.

2.3 A worked established example: mixed Tate motives

Example 2.6 (Mixed Tate motives as a neutral Tannakian category). [EST] For a number field k and a finite set S of places, the category $\text{MTM}(k; S)$ of mixed Tate motives unramified outside S is an unconditionally constructed neutral Tannakian category [9, 10]. Its fiber functors include the Betti, de Rham, Hodge, and ℓ -adic étale realizations; its Tannakian dual is the *motivic Galois group*, an extension of $\mathbb{G}_m$ by a prounipotent group whose Lie algebra is a free graded Lie algebra with one generator in each odd degree ≥ 3 . The period pairing between the Betti and de Rham realizations produces the numbers $\zeta(n)$ and multiple zeta values as *periods*, i.e. as *observables* $\text{per}(\text{motive})$ [26]. This is a fully rigorous instance of the ladder (1) with target $\text{Vect}_{\mathbb{Q}}^{\text{fd}}$: abstract algebra (MTM), realization (a fiber functor), observable (a period). Part V proposes to enlarge the target from $\text{Vect}_{\mathbb{Q}}$ to $\text{Cond}(\text{Vect})$; Example 2.6 is the concrete precedent that makes the proposal legible.

2.4 The finite case, in full

To make Theorem 2.4 completely explicit (and to underwrite the accompanying code), we record the finite-group case with a self-contained proof.

Proposition 2.7 (Reconstruction of a finite group). [EST] *Let G be a finite group and k an algebraically closed field with $\text{char}(k) \nmid |G|$. Let $\mathcal{C} = \text{Rep}_k(G)$ be the category of finite-dimensional k -representations, a semisimple rigid symmetric monoidal k -linear category, and let $\omega: \text{Rep}_k(G) \rightarrow \text{Vect}_k^{\text{fd}}$ be the forgetful functor. Then the group of monoidal natural automorphisms of ω satisfies*

$$\text{Aut}^{\otimes}(\omega) \cong G,$$

naturally; equivalently, G is recovered as the group of k -points of the finite constant group scheme $\underline{\text{Aut}}^{\otimes}(\omega)$.

Proof. A monoidal natural transformation $\lambda: \omega \Rightarrow \omega$ assigns to each representation (V, ρ_V) a linear automorphism $\lambda_V \in \text{GL}(V)$ such that (naturality) for every G -map $f: V \rightarrow W$, $f \circ \lambda_V = \lambda_W \circ f$, and (monoidality) $\lambda_{V \otimes W} = \lambda_V \otimes \lambda_W$ and $\lambda_{\mathbf{1}} = \text{id}$.

Each element of G gives such a λ . For $g \in G$ set $\lambda_V := \rho_V(g)$. Naturality is G -equivariance of f ; monoidality is $\rho_{V \otimes W}(g) = \rho_V(g) \otimes \rho_W(g)$ and $\rho_{\mathbf{1}}(g) = \text{id}$. Distinct elements give distinct transformations because the regular representation $k[G]$ is faithful: if $\rho_{k[G]}(g) = \rho_{k[G]}(h)$ then $g = h$. Hence $G \hookrightarrow \text{Aut}^{\otimes}(\omega)$.

Every such λ comes from G . Let $R := k[G]$ be the left regular representation, with G acting by left translations $\rho_R(g) = L_g$. The G -endomorphisms of R are exactly the right translations, $\text{End}_G(R) = \{R_g : g \in G\}$ (the commutant of the left regular representation is the right regular representation). Naturality of λ at each $f \in \text{End}_G(R)$ forces $\lambda_R f = f \lambda_R$; hence λ_R lies in the commutant of $\{R_g\}$, which is the linear span of the left translations, $\lambda_R = \sum_g c_g L_g$ with $c_g \in k$.

Now we use the *coalgebra* structure, not the algebra structure. The comultiplication $\Delta: R \rightarrow R \otimes R$, $\Delta(h) = h \otimes h$, is a morphism in $\text{Rep}_k(G)$: the basis of R consists of group elements and G acts diagonally on $R \otimes R$, so under the left regular action on R one has $\Delta(L_g h) = gh \otimes gh = (L_g \otimes L_g) \Delta(h)$. (By contrast the multiplication $m: R \otimes R \rightarrow R$ is not a G -map for non-abelian G — $m(g \cdot (x \otimes y)) = gxgy \neq gxy = g \cdot m(x \otimes y)$ — so m

cannot be used here.) Naturality of λ at Δ , together with monoidality $\lambda_{R \otimes R} = \lambda_R \otimes \lambda_R$, gives $(\lambda_R \otimes \lambda_R) \circ \Delta = \Delta \circ \lambda_R$, so λ_R is a coalgebra endomorphism. Evaluating on a basis element h ,

$$\sum_g c_g (gh \otimes gh) = \Delta(\lambda_R h) = (\lambda_R \otimes \lambda_R) \Delta(h) = \sum_{g, g'} c_g c_{g'} (gh \otimes g'h),$$

and comparing coefficients on the basis $\{gh \otimes g'h\}$ forces $c_g c_{g'} = 0$ for $g \neq g'$ and $c_g^2 = c_g$. Thus each $c_g \in \{0, 1\}$ with at most one nonzero; invertibility of λ_R selects exactly one, so $\lambda_R = L_{g_0}$ for a unique $g_0 \in G$. Finally every finite-dimensional V is a subrepresentation of a finite sum $R^{\oplus m}$ (each irreducible occurs in R by Maschke's theorem, using $\text{char}(k) \nmid |G|$), and the inclusion and projections are G -maps; naturality at these propagates $\lambda_V = \rho_V(g_0)$ to every V . Hence $\text{Aut}^\otimes(\omega) \hookrightarrow G$, and the two inclusions are mutually inverse. $\square$

Proposition 2.7 is the smallest complete instance of “geometry from algebra”: the finite group G — a set with structure, the crudest “geometric” datum — is recovered exactly from the tensor category $\text{Rep}_k(G)$ and the fiber functor ω . Our code (Section 10) carries this out for concrete small groups by computing the tensor-automorphisms directly from character/fusion data.

3 Four more precedents, one schema

We now place the remaining established reconstruction theorems beside Tannakian duality and extract their common form. Each recovers a geometric or dynamical object from algebraic data; each is a theorem, not a heuristic.

3.1 Gelfand duality: a space from a commutative C^* -algebra

Theorem 3.1 (Gelfand–Naimark). [EST] *The functor $A \mapsto \text{Spec}(A) = \text{Hom}_{C^*}(A, \mathbb{C})$ (characters, with the weak- $*$ topology) is a contravariant equivalence between the category of commutative unital C^* -algebras and the category of compact Hausdorff spaces, with quasi-inverse $X \mapsto C(X)$; the Gelfand transform $A \rightarrow C(\text{Spec} A)$ is an isometric $*$ -isomorphism [11].*

Gelfand duality is the archetype “space from algebra”: the topological space X is literally the set of algebra homomorphisms $A \rightarrow \mathbb{C}$, reconstructed with its topology, functorially and naturally. It is the commutative shadow of noncommutative geometry — a spectral triple $(\mathcal{A}, \mathcal{H}, D)$ with $\mathcal{A}$ commutative has $\text{Spec}(\mathcal{A})$ a space, and the extra datum D upgrades “space” to “metric space” (Section 5). We regard Theorem 3.1 as the *objectwise* half of the ladder and Tannakian duality as the *tensor/symmetry* half.

Remark 3.2 (Naturality is the load-bearing word). In all these theorems, “recovered” means *recovered by a natural equivalence of categories*, not merely “in bijection with.” Naturality is what makes the reconstruction a *functor* and forbids arbitrary ad hoc choices; it is also what the collaboration’s Equivalence axiom demands (Remark 2.5). Whenever we later speak of an emergent geometry we mean an assignment that is functorial in the algebraic input, so that a morphism of algebraic data induces a morphism of geometries.

3.2 Connes' spectral reconstruction: a Riemannian manifold from a spectral triple

Definition 3.3 (Spectral triple). A *spectral triple* $(\mathcal{A}, \mathcal{H}, D)$ consists of a $*$ -algebra $\mathcal{A}$ represented on a Hilbert space $\mathcal{H}$, together with a self-adjoint operator D (generally unbounded) with compact resolvent $(D + i)^{-1}$ and bounded commutators $[D, a]$ for all $a \in \mathcal{A}$. It is *even* if there is a $\mathbb{Z}/2$ -grading γ with $\gamma a = a\gamma$, $\gamma D = -D\gamma$; it is *real* of KO-dimension n if there is an antiunitary J implementing the requisite commutation relations.

Theorem 3.4 (Connes' spectral distance formula). [EST] *For the canonical spectral triple $(C^\infty(M), L^2(M, \mathcal{S}), \not{D})$ of a closed Riemannian spin manifold M with Dirac operator $\not{D}$, the geodesic distance is*

$$d_{\not{D}}(p, q) = \sup\{|f(p) - f(q)| : f \in C^\infty(M), \|\llbracket \not{D}, f \rrbracket\| \leq 1\}, \quad p, q \in M.$$

The condition $\|\llbracket \not{D}, f \rrbracket\| \leq 1$ is equivalent to f being 1-Lipschitz, and the supremum recovers the Riemannian metric purely from the operator-algebraic pair $(C^\infty(M), \not{D})$ [13].

Theorem 3.5 (Connes' reconstruction theorem). [EST] *Let $(\mathcal{A}, \mathcal{H}, D)$ be a commutative real spectral triple of KO-dimension n satisfying Connes' axioms: dimension/regularity, finiteness and absolute continuity, the first-order condition, orientability (a Hochschild cycle representing the volume form), Poincaré duality in K -theory, and the reality/grading conditions. Then there is a closed oriented Riemannian spin manifold M of dimension n with $\mathcal{A} \cong C^\infty(M)$ and $(\mathcal{A}, \mathcal{H}, D)$ unitarily equivalent to the canonical spectral triple of M [12].*

Theorem 3.5 is, of all the precedents, the closest to the physical ambition of this Part: *a manifold with its metric is recovered as an invariant of purely spectral/algebraic data, with no manifold assumed in the input.* It is the rigorous kernel around which the SPEC proposal of Section 6 is organized. (Its dynamical extension, the Chamseddine–Connes spectral action $\text{tr } f(D/\Lambda)$ [22], reproduces the Einstein–Hilbert action from the same spectral data via a heat-kernel expansion, but at only [HEU] status — a template, not a theorem, for how dynamics might follow the reconstruction.) We flag immediately, and will repeat, its decisive limitation (Section 6.3): it is *Riemannian*.

3.3 Tomita–Takesaki and Bisognano–Wichmann: boosts from a von Neumann algebra

Theorem 3.6 (Tomita–Takesaki). [EST] *Let $\mathfrak{M} \subseteq B(\mathcal{H})$ be a von Neumann algebra with a cyclic and separating vector Ω . Let S be the closure of $a\Omega \mapsto a^*\Omega$, with polar decomposition $S = J\Delta^{1/2}$. Then J is an antiunitary involution, $\Delta > 0$ is self-adjoint, $\Delta^{it}\mathfrak{M}\Delta^{-it} = \mathfrak{M}$ for all $t \in \mathbb{R}$ (the modular flow σ_t), and $J\mathfrak{M}J = \mathfrak{M}'$ [15].*

Theorem 3.7 (Bisognano–Wichmann). [EST] *For a Wightman quantum field theory, let $\mathfrak{M} = \mathcal{A}(W)$ be the von Neumann algebra of observables localized in the Rindler wedge $W = \{x^1 > |x^0|\}$ and let Ω be the vacuum. Then the modular flow is the boost, $\sigma_t = \text{Ad } U(\Lambda_W(-2\pi t))$, where $\Lambda_W(s)$ is the one-parameter group of Lorentz boosts preserving W ; and $J = \Theta U(R_1(\pi))$ combines the antiunitary PCT operator Θ (the product of parity, charge*

conjugation, and time reversal) with the unitary $U(R_1(\pi))$ implementing a spatial rotation by angle π about the x^1 -axis. Consequently the modular flow is geometric — a Lorentz boost — and Haag duality $\mathcal{A}(W)' = \mathcal{A}(W')$ holds for the causal complement W' [16, 17].

Theorem 3.7 is the one established precedent that recovers *Lorentzian* data — a boost, its associated Unruh temperature $T = 1/2\pi$, and the causal complement — from purely operator-algebraic input $(\mathfrak{M}, \Omega)$. It is therefore the crucial pointer toward what a Lorentzian reconstruction should look like, and simultaneously a caution: Bisognano–Wichmann *pre-supposes* Minkowski space and Poincaré covariance in its hypotheses; it recovers the boost as *modular*, but it does not build spacetime from scratch. We return to this tension in Section 6.3.

3.4 The cobordism hypothesis: physics as a rigidly determined functor

Theorem 3.8 (Cobordism hypothesis; Baez–Dolan, Lurie). [EST] *Let $\mathcal{C}$ be a symmetric monoidal (∞, n) -category with duals. Evaluation on the point, $Z \mapsto Z(\text{pt})$, is an equivalence between the ∞ -groupoid of symmetric monoidal functors $\mathbf{Bord}_n^{\text{fr}} \rightarrow \mathcal{C}$ (framed fully extended TQFTs) and the ∞ -groupoid $\mathcal{C}^\sim$ of fully dualizable objects of $\mathcal{C}$ [18].*

The cobordism hypothesis is the purest form of “physics as a functor”: a field theory $Z: \mathbf{Bord}_n \rightarrow \mathcal{C}$ is a monoidal functor from a geometric category (cobordisms) to an algebraic one, and it is reconstructed in full from a single object $Z(\text{pt})$. This is the direct ancestor, in the prior-work corpus, of Part V’s “realization functor” language: a TQFT literally is a realization functor from geometry to algebra; the cobordism hypothesis says such functors are rigid. Part V contemplates the *reverse* arrow — from algebra to geometry — and uses Theorem 3.8 as evidence that “physics as a functor between a geometric and an algebraic category” is a mathematically robust, not merely rhetorical, idea.

3.5 The common schema

Principle 3.9 (Reconstruction schema). [HEU] Each of Theorems 2.4, 3.1, 3.5, 3.7 and 3.8 instantiates:

There is a functor Real from a category $\mathbf{Alg}$ of algebraic/representation data to a category $\mathbf{Geom}$ of geometric/dynamical objects, and a reconstruction functor $\text{Rec}: \mathbf{Geom} \rightarrow \mathbf{Alg}$, such that Rec is fully faithful and Real is (an inverse or a section) on its essential image; the geometric object is recovered, up to natural equivalence, as $\text{Real}(\text{its algebra})$.

Theorem	Algebraic input	Geometric output	Signature
Gelfand	commutative C^* -algebra A	compact Hausdorff $\mathrm{Spec} A$	topology
Tannaka	tensor category + fiber functor	affine group scheme G_ω	scheme
Connes	commutative spectral triple $(\mathcal{A}, \mathcal{H}, D)$	Riemannian spin manifold	metric (Riem.)
Bisognano– Wichmann	von Neumann algebra $(\mathfrak{M}, \Omega)$	boost, Unruh T , causal compl.	Lorentzian (local)
Cobordism hyp.	fully dualizable object $Z(\mathrm{pt})$	TQFT functor Z on Bord_n	topological

Principle 3.9 is [HEU], not [EST]: it is a structural reading uniting five theorems, not itself a theorem. What is [EST] is each row. The physical program of this Part is to add a sixth row — with *condensed* algebraic input and *emergent (Lorentzian) spacetime* output — and the honest status of that row is [SPEC].

4 The realization functor and the representation ladder

We now make the ladder (1) precise and state the Realization axiom in condensed form.

4.1 Realization functors, abstractly

Definition 4.1 (Realization functor). Let $\mathcal{C}$ be a rigid k -linear tensor category (the *abstract algebra*) and let $\mathcal{T}$ be a rigid symmetric monoidal k -linear category (the *target of realization*). A *realization functor* is an exact, faithful, k -linear symmetric monoidal functor

$$\omega: \mathcal{C} \longrightarrow \mathcal{T}.$$

When $\mathcal{T} = \mathrm{Vect}_k^{\mathrm{fd}}$ this is a fiber functor and (if $\mathcal{C}$ is neutral Tannakian) Theorem 2.4 applies. We are interested in enlarging $\mathcal{T}$ to a category of *condensed* vector spaces, $\mathcal{T} = \mathrm{Cond}(\mathrm{Vect}_k)$ or a solid/liquid subcategory thereof.

Definition 4.2 (Condensed fiber functor). A *condensed fiber functor* on a k -linear symmetric monoidal abelian category $\mathcal{C}$ with $\mathrm{End}(\mathbf{1}) = k$ is an exact, faithful, k -linear symmetric monoidal functor $\omega: \mathcal{C} \rightarrow \mathrm{Cond}(\mathrm{Vect}_k)$ landing in the *reflexive* condensed vector spaces — those M for which the canonical map $M \rightarrow M^{\vee\vee}$ is an isomorphism, e.g. suitable solid or liquid modules. We call $(\mathcal{C}, \omega)$ a *condensed Tannakian category* (in this reflexive, not strictly rigid, sense).

Remark 4.3 (The rigidity obstruction — why reflexivity, not rigidity). [EST] A genuine obstruction forces the weakening in Definition 4.2, and it must be flagged. In a complete symmetric monoidal category of topological or condensed vector spaces, the strictly *dualizable* (rigid) objects are exactly the finite-dimensional discrete ones: the coevaluation $\mathbf{1} \rightarrow M \otimes M^\vee$ must factor through finite-rank tensors, so an infinite-dimensional solid or liquid module is *not* dualizable. Consequently, if one demanded that $\mathcal{C}$ be rigid and that ω preserve duals, the image of ω would be forced finite-dimensional and the entire “condensed” enrichment would collapse back to classical Tannakian reconstruction (Theorem 2.4), defeating the purpose of enlarging the target (Remark 4.4). Definition 4.2 therefore trades strict rigidity for reflexivity, at a real cost: the classical reconstruction machinery does *not* transfer verbatim, and a reconstruction theorem in the reflexive/non-rigid setting must be proved anew. This obstruction is precisely (part of) frontier (C) in Remark 6.3, and it is one concrete reason the condensed Tannakian program is a nontrivial open problem rather than a routine generalization.

Definition 4.2 is well-posed as a definition: $\text{Cond}(\text{Vect}_k)$ is a k -linear symmetric monoidal category (indeed the solid and liquid subcategories are, by Clausen–Scholze, well-behaved symmetric monoidal abelian/stable categories [20, 21]), so “exact, faithful, monoidal functor into it” has content. What is *not* established — and is the SPEC proposal of Section 6 — is any reconstruction theorem asserting that such an ω reconstructs a geometry. Definition 4.2 is the honest bookkeeping of what a “condensed fiber functor” *is*; Section 6 states what we *conjecture* it does.

Remark 4.4 (Why enlarge the target). The classical target $\text{Vect}_k^{\text{fd}}$ has no room for topology, analysis, or the “continuous families indexed by profinite sets” that carry physical data (Part IV’s fields, Part II’s observables). Condensed vector spaces are exactly the enlargement that (i) contains ordinary topological vector spaces faithfully (on compactly generated objects), (ii) forms a genuine abelian category supporting homological algebra, and (iii) carries a symmetric monoidal structure (solid/liquid tensor) with internal Hom. Enlarging the fiber-functor target from Vect_k to $\text{Cond}(\text{Vect}_k)$ is therefore the minimal move that could let a Tannakian-style reconstruction “see” geometric/analytic structure rather than only a bare group scheme. This is the technical rationale for the ladder (1).

4.2 The ladder as a commuting diagram

The five-stage ladder (1) refines to the following diagram, in which the top row is the *abstract* pipeline and the bottom row is a *realized* instance (a choice of fiber functor ω and a state ϕ):

$$\begin{array}{ccccccc}
 \mathcal{C} & \xrightarrow{\iota} & \mathcal{C}^{\text{Cond}} & \xrightarrow{\omega^{\text{Cond}}} & \text{Cond}(\text{Vect}_k) & \xrightarrow{\phi} & \text{states} \xrightarrow{d} (X, g) \\
 \omega \downarrow & & \downarrow & & (-)(*) \downarrow & & \downarrow \\
 \text{Vect}_k^{\text{fd}} & \xrightarrow{\iota} & \text{Cond}(\text{Vect}_k) & = & \text{Cond}(\text{Vect}_k) & \xrightarrow{\phi} & \mathbb{R} \text{ or } \mathbb{C} \xrightarrow{d} (X, g)
 \end{array}$$

(Note: The bottom row in the original image has an additional arrow from $\text{Cond}(\text{Vect}_k)$ to $\mathbb{R} \text{ or } \mathbb{C}$ labeled ϕ , and an arrow from states to (X, g) labeled d . The vertical arrow from states to $\mathbb{R} \text{ or } \mathbb{C}$ is labeled $\cong$.)

Reading left to right: an abstract tensor category $\mathcal{C}$ is embedded (ι) into a condensed enrichment $\mathcal{C}^{\text{Cond}}$; a condensed fiber functor ω^{Cond} realizes objects as condensed vector spaces; evaluation on a state ϕ (a condensed linear functional, in the sense of Part IV) produces

numerical observables; and a distance/metric read-off d (in the style of Connes, Section 5) assembles the observables into a geometry (X, g) . The classical Tannakian and Connes theorems are the assertions that specific sub-arrows of this diagram are equivalences; the condensed enrichment of the whole is the open problem.

Definition 4.5 (Condensed Realization axiom). [HEU] The *condensed Realization axiom* asserts that physical content is the image of an abstract condensed object under a realization functor: for a condensed tensor category $\mathcal{C}^{\text{Cond}}$ of observables there is a realization $\omega^{\text{Cond}}: \mathcal{C}^{\text{Cond}} \rightarrow \text{Cond}(\text{Vect}_k)$ and a family of states such that all measurable predictions are values $\phi(\omega^{\text{Cond}}(X))$, and any two realizations differ by a torsor under the reconstructed symmetry (the condensed Tannakian dual).

Definition 4.5 is [HEU]: it is the condensed instance of the collaboration’s Realization and Equivalence axioms [25], physically motivated and structurally precise, but not a theorem. Its EST kernel is Theorem 2.4 (the classical torsor statement); its SPEC extension is that the reconstructed symmetry organizes an *emergent geometry*, treated in Section 6.

5 Emergent metrics: spectral distance as the EST toy model

The fifth arrow d of the ladder — “read a metric off the realization” — is the one for which there is a rigorous, computable, established prototype: Connes’ spectral distance. We develop it in the finite setting, where it is completely elementary, entirely rigorous, and reproduced by our code.

5.1 Finite spectral triples and the spectral distance

Definition 5.1 (Finite spectral triple). A *finite spectral triple* is a triple $(\mathcal{A}, \mathcal{H}, D)$ with $\mathcal{H}$ a finite-dimensional Hilbert space, $\mathcal{A} \subseteq B(\mathcal{H})$ a finite-dimensional commutative $*$ -algebra with unit represented on $\mathcal{H}$, and $D = D^* \in B(\mathcal{H})$. Since $\mathcal{A}$ is commutative and finite-dimensional it is $*$ -isomorphic to $\mathbb{C}^N$ for $N = \dim_{\mathbb{C}} \mathcal{A}$; its *spectrum* $\text{Spec} \mathcal{A} = \{p_1, \dots, p_N\}$ is a finite set (the pure states / characters), and each $a \in \mathcal{A}$ is a function $a: \text{Spec} \mathcal{A} \rightarrow \mathbb{C}$, $p_i \mapsto a(p_i)$.

Definition 5.2 (Spectral distance on a finite triple). For a finite spectral triple $(\mathcal{A}, \mathcal{H}, D)$ the *spectral (Connes) distance* between pure states $p, q \in \text{Spec} \mathcal{A}$ is

$$d_D(p, q) = \sup \{ |a(p) - a(q)| : a \in \mathcal{A}, \| [D, a] \| \leq 1 \} \in [0, \infty].$$

Proposition 5.3 (The spectral distance is an extended metric). [EST] For any finite spectral triple, d_D is symmetric, satisfies the triangle inequality, and $d_D(p, q) = 0$ iff $p = q$ (provided the $[D, \cdot]$ -seminorm separates points, i.e. D does not commute with all of $\mathcal{A}$ across any pair of points). Hence d_D is an extended metric on $\text{Spec} \mathcal{A}$ [13].

Proof. Symmetry is immediate (a and $-a$ have equal seminorm). For the triangle inequality, if a is admissible ($\| [D, a] \| \leq 1$) then $|a(p) - a(r)| \leq |a(p) - a(q)| + |a(q) - a(r)| \leq d_D(p, q) + d_D(q, r)$, and taking the sup over admissible a gives $d_D(p, r) \leq d_D(p, q) + d_D(q, r)$. If $p = q$

then $|a(p) - a(q)| = 0$ for all a , so $d_D(p, q) = 0$. Conversely if $p \neq q$, choose a with $a(p) \neq a(q)$ and $\|[D, a]\| > 0$ (possible when D links p, q); rescaling a to unit seminorm gives $|a(p) - a(q)| > 0$, so $d_D(p, q) > 0$. $\square$

5.2 The two-point space: a metric from a single off-diagonal entry

Example 5.4 (Connes' two-point space). [EST] Let $\mathcal{A} = \mathbb{C}^2 = \{(a_1, a_2)\}$ act diagonally on $\mathcal{H} = \mathbb{C}^2$, and let $D = \begin{pmatrix} 0 & \bar{m} \\ m & 0 \end{pmatrix}$ with $m \in \mathbb{C}^\times$. For $a = (a_1, a_2)$, $[D, a] = \begin{pmatrix} 0 & \bar{m}(a_2 - a_1) \\ m(a_1 - a_2) & 0 \end{pmatrix}$, whose operator norm is $|m| |a_1 - a_2|$. The constraint $\|[D, a]\| \leq 1$ reads $|a_1 - a_2| \leq 1/|m|$, so

$$d_D(p_1, p_2) = \sup\{|a_1 - a_2| : |m| |a_1 - a_2| \leq 1\} = \frac{1}{|m|}.$$

The distance between the two points is the inverse of the off-diagonal Dirac coupling. This is the smallest nontrivial instance of “metric from spectral data”: the geometry (a two-point metric space of diameter $1/|m|$) is read off from a single algebraic number m . Our code reproduces this exactly.

5.3 Graph metrics as spectral distances

The finite spectral distance recovers, for a natural class of triples, the *shortest-path (graph) metric* — a fully combinatorial, checkable “emergent geometry.”

Construction 5.5 (Weighted-graph triple as a direct sum of two-point spaces). Let $\Gamma = (V, E, w)$ be a finite connected graph with $N = |V| \geq 2$ vertices $V = \{1, \dots, N\}$ and positive edge weights $w_{ij} > 0$ for $\{i, j\} \in E$. (Connectedness guarantees that every vertex meets some edge block, so the representation of $\mathcal{A} = \mathbb{C}^V$ on $\mathcal{H}$ is faithful and no vertex sits at infinite distance.) We build the triple as the *direct sum, over edges, of the two-point spectral triples of Example 5.4*, so that the commutator is block-diagonal and its operator norm is a maximum over edges rather than a degree-coupled quantity. Concretely, put one two-dimensional block per edge,

$$\mathcal{H} = \bigoplus_{e=\{i,j\} \in E} \mathbb{C}_e^2, \quad \mathcal{A} = \mathbb{C}^V \text{ acting on the block } \mathbb{C}_e^2 \text{ (} e = \{i, j\} \text{) by } a \mapsto \text{diag}(a_i, a_j),$$

and let $D = \bigoplus_e D_e$ act blockwise by the two-point Dirac operator of coupling w_{ij}^{-1} ,

$$D_e = \begin{pmatrix} 0 & w_{ij}^{-1} \\ w_{ij}^{-1} & 0 \end{pmatrix} \text{ on } \mathbb{C}_e^2, \quad e = \{i, j\}.$$

This is manifestly a finite spectral triple (Definition 5.1); the earlier, naive “vertex-to-edge coupling” Dirac operator is deliberately avoided because its commutator norm is governed by the graph Laplacian spectrum and scales with vertex degree, which would make the recovered distance strictly shorter than the geodesic.

Theorem 5.6 (Spectral distance recovers the graph geodesic). [EST] *For the finite spectral triple of Construction 5.5, the admissibility constraint $\|[D, a]\| \leq 1$ implies $|a_i - a_j| \leq w_{ij}$ for every edge $\{i, j\}$ (each edge is a “1-Lipschitz” constraint of width w_{ij}), and conversely any a satisfying all edge constraints is admissible. Consequently the spectral distance equals the weighted shortest-path distance:*

$$d_D(p, q) = d_\Gamma(p, q) := \min_{\text{paths } p=v_0, \dots, v_\ell=q} \sum_{s=1}^{\ell} w_{v_{s-1}v_s}.$$

Proof sketch. By Construction 5.5 the commutator is block-diagonal,

$$[D, a] = \bigoplus_{e=\{i,j\} \in E} [D_e, \text{diag}(a_i, a_j)],$$

so its operator norm is the maximum of the block norms; by the two-point computation (Example 5.4) the block for $e = \{i, j\}$ has norm $w_{ij}^{-1}|a_i - a_j|$. Hence $\|[D, a]\| = \max_{\{i,j\} \in E} w_{ij}^{-1}|a_i - a_j|$, and $\|[D, a]\| \leq 1 \iff |a_i - a_j| \leq w_{ij}$ for every edge. The optimization

$$d_D(p, q) = \sup\{a(p) - a(q) : |a_i - a_j| \leq w_{ij} \ \forall \{i, j\} \in E\}$$

is a linear program whose feasible potentials a are exactly the 1-Lipschitz functions for the path metric d_Γ . Its optimum is $d_\Gamma(p, q)$: the function $a_i := d_\Gamma(p, i)$ is feasible (the triangle inequality gives $|d_\Gamma(p, i) - d_\Gamma(p, j)| \leq w_{ij}$) and achieves $a(p) - a(q) = -d_\Gamma(p, q)$ up to sign, while no feasible a can exceed $d_\Gamma(p, q)$ by telescoping along any shortest path. This is the Kantorovich–Rubinstein duality specialized to a finite metric graph. $\square$

Theorem 5.6 is the cleanest available “emergent geometry” statement: a finite metric space (the weighted graph geodesic) is recovered as a spectral distance from an algebra and a discrete Dirac operator. It is entirely [EST], it is a finite linear program, and our code computes it for explicit graphs and checks $d_D = d_\Gamma$ against a Floyd–Warshall shortest-path reference. This is the toy model that makes the fifth arrow d of the ladder (1) concrete.

Heuristic 5.7 (Metric from realization data). [HEU] Theorems 3.4 and 5.6 and Example 5.4 together support the dictionary entry: *given a realization of algebraic data as operators on a Hilbert space and a distinguished “Dirac” operator D , a metric is read off as the D -Lipschitz-dual distance on the state space.* As a heuristic this is robust (it is a theorem in each cited case); as a route to *physical spacetime* it inherits the Riemannian limitation of its established instances (Section 6.3).

6 The speculative core: condensed reconstruction of spacetime

We now state the central hypothesis of this Part, box it as [SPEC], and dissect the gap that separates it from the established precedents.

6.1 The hypothesis

SPEC — Spacetime as the realization of condensed data

Hypothesis 6.1 (Condensed realization of spacetime). [SPEC] There is a condensed Tannakian category $(\mathcal{C}^{\text{Cond}}, \omega^{\text{Cond}})$ (Definition 4.2) — a rigid k -linear condensed tensor category of “observables” with a condensed fiber functor into $\text{Cond}(\text{Vect}_k)$ — and a reconstruction functor

$$\text{Rec}: \{\text{emergent geometries}\} \longrightarrow \{\text{condensed Tannakian data}\},$$

such that physical spacetime (X, g) is recovered, up to natural equivalence, as $(X, g) \cong \text{Real}(\mathcal{C}^{\text{Cond}}, \omega^{\text{Cond}})$, with the metric read off by a condensed spectral distance (Heuristic 5.7) and the causal/Lorentzian structure recovered by a Tomita–Takesaki–type modular flow (Theorem 3.7). Equivalently: *spacetime is the realization of condensed algebraic data.*

Hypothesis 6.1 is not a theorem and is not claimed to be one. It is the precise, boxed articulation of the Part’s title. Its plausibility rests on the five established precedents; its unproven content is *everything condensed and everything Lorentzian*. We now separate those two frontiers.

6.2 What is established, what is proposed

Proposition 6.2 (The established sub-ladder). [EST] *The following sub-arrows of Hypothesis 6.1 are established:*

1. (Tannaka) For $\mathcal{T} = \text{Vect}_k^{\text{fd}}$, a neutral Tannakian category reconstructs an affine group scheme (Theorem 2.4); the target enlargement to $\text{Cond}(\text{Vect}_k)$ leaves the input side $(\mathcal{C}, \omega)$ a well-defined structure (Definition 4.2).
2. (Metric) A metric is recoverable from a spectral triple by the spectral distance in the Riemannian and finite cases (Theorems 3.4 and 5.6).
3. (Modular) Boosts, Unruh temperature, and the causal complement are recoverable from a von Neumann algebra with a cyclic–separating vector (Theorem 3.7).
4. (Functoriality) “Physics as a functor” from a geometric to an algebraic category is mathematically robust (Theorem 3.8).

Proof. Each clause is a restatement of the cited established theorem; clause (1)’s enlargement claim is Definition 4.2 together with the fact that solid/liquid modules form symmetric monoidal abelian categories [20, 21]. □

SPEC — the unproven bridges

Remark 6.3 (The two frontiers). [SPEC] Hypothesis 6.1 requires two bridges that no cited theorem supplies:

- (C) *The condensed frontier.* No reconstruction theorem is known for condensed fiber functors: given $(\mathcal{C}^{\text{Cond}}, \omega^{\text{Cond}})$ with $\mathcal{T} = \text{Cond}(\text{Vect}_k)$, there is no established statement recovering a geometry (as opposed to a bare group scheme). Moreover the rigidity obstruction (Remark 4.3) shows the categorical setup itself must be weakened — reflexive rather than rigid objects — so the classical machinery does not transfer verbatim. Constructing a *condensed Tannakian reconstruction theorem* in the reflexive setting is an open problem.
- (L) *The Lorentzian frontier.* Every established metric reconstruction (Theorems 3.5 and 5.6) is *Riemannian*. A reconstruction recovering *Lorentzian causal structure* (not just distance) from spectral/representation data is open [14]. Bisognano–Wichmann recovers a boost, but from a theory already living on Minkowski space; it does not build the causal structure ex nihilo.

Frontiers (C) and (L) are logically independent; Hypothesis 6.1 asserts both, and by the worst-component-wins discipline its status is **[SPEC]**.

6.3 The Riemannian–Lorentzian gap, honestly

We isolate frontier (L) because it is the sharpest and most often glossed. The issue is not merely a sign: a Lorentzian “metric” has an indefinite signature, the Dirac operator becomes hyperbolic rather than elliptic, its resolvent is no longer compact, and the spectral-distance supremum degenerates (spacelike-separated points are at “distance zero” or the sup is $+\infty$). Concretely:

- In the Riemannian case, $\not{D}$ is elliptic with compact resolvent; the spectral distance is finite and metrizes the topology (Theorem 3.4). This is what Connes’ theorem uses.
- In the Lorentzian case, the natural Dirac operator is hyperbolic; the “distance” induced by a Lipschitz-type supremum is a *Lorentzian distance* (longest causal path, reverse triangle inequality), and recovering it requires selecting a causal (time) order, which is extra data not present in a bare spectral triple. Franco–Eckstein [14] formulate an algebraic causality condition (a cone of “causal functions”) but do not yield a full reconstruction theorem.

Principle 6.4 (Non-identification: Riemannian reconstruction $\neq$ Lorentzian reconstruction). **[EST]** (as a statement about the literature). Connes’ reconstruction theorem (Theorem 3.5) is proved in Riemannian signature and recovers a metric *space*; it does *not* establish, and must not be quoted as establishing, the recovery of a Lorentzian spacetime with its causal structure. The Lorentzian reconstruction problem is open. Any composite claim of this program that invokes an emergent *Lorentzian* geometry is therefore **[SPEC]**, regardless of how rigorous the Riemannian ingredients are.

Principle 6.4 is the single most important honesty constraint of this Part. We state it as a labeled principle precisely so that no downstream synthesis (Parts VI–VII) can silently upgrade the Riemannian EST results into a Lorentzian claim.

6.4 A concrete research target

To keep Hypothesis 6.1 from being merely a slogan, we isolate a sharp, potentially attackable sub-question.

Conjecture 6.5 (Condensed spectral distance recovers a condensed metric space). [SPEC] Let $(\mathcal{A}^{\text{Cond}}, \mathcal{H}^{\text{Cond}}, D)$ be a “condensed spectral triple”: $\mathcal{A}^{\text{Cond}}$ a commutative algebra object in $\text{Cond}(\text{Vect}_{\mathbb{R}})$ (e.g. a solid $\mathbb{R}$ -algebra), $\mathcal{H}^{\text{Cond}}$ a condensed Hilbert module, D a self-adjoint condensed operator with bounded commutators. Then the condensed spectral distance $d_D(p, q) = \sup\{|p(a) - q(a)| : \|[D, a]\| \leq 1\}$, taken over condensed characters p, q , defines a metric on the condensed spectrum $\text{Spec}_{\text{Cond}}(\mathcal{A}^{\text{Cond}})$ that refines the Gelfand topology, and, when $\mathcal{A}^{\text{Cond}}$ is the condensed avatar of $C^\infty(M)$ for a compact Riemannian M , recovers the geodesic metric of M .

Conjecture 6.5 is deliberately *Riemannian and static* — it isolates frontier (C) while holding frontier (L) fixed — so that it is a well-posed mathematical target rather than a physical aspiration. It is, to our knowledge, open; the finite version (Theorem 5.6) is its established shadow, and the condensed avatar of the compact-manifold case is the natural first theorem to prove. Frontier (L) — upgrading Conjecture 6.5 to Lorentzian signature and dynamical (Einstein-type) content — is deferred to Parts VI–VII and remains [SPEC].

7 Results

We summarize the paper’s deliverables, tagged by status.

1. **A unified schema** ([HEU]). The reconstruction schema Principle 3.9 organizes five established theorems — Gelfand, Tannaka, Connes, Bisognano–Wichmann, cobordism hypothesis — as instances of “geometry is the realization of algebraic data,” with an explicit signature table.
2. **The representation ladder** ([HEU]/[EST]). Definitions 4.1 and 4.2 and the commuting diagram of Section 4 make the ladder $\mathcal{C} \rightarrow \mathcal{C}^{\text{Cond}} \rightarrow \text{Cond}(\text{Vect}) \rightarrow \text{states} \rightarrow (X, g)$ precise; its left-hand sub-arrows are [EST] (Theorem 2.4), its condensed enrichment is the open problem.
3. **Finite Tannakian reconstruction** ([EST]). Proposition 2.7 gives a complete elementary proof that $\text{Aut}^\otimes(\omega) \cong G$ for a finite group; the accompanying code recovers concrete small groups from their representation/fusion data.
4. **Spectral distance as emergent metric** ([EST]). Theorem 5.6 proves that the finite Connes spectral distance recovers the weighted graph geodesic (a finite linear program / Kantorovich–Rubinstein duality), with Example 5.4 the two-point base case; the code computes d_D and checks it against Floyd–Warshall.
5. **The boxed SPEC hypothesis and its dissection** ([SPEC]). Hypothesis 6.1 states “spacetime is the realization of condensed data”; Remark 6.3 separates the condensed

frontier (C) from the Lorentzian frontier (L); Conjecture 6.5 isolates a well-posed Riemannian condensed target; Principle 6.4 records the non-negotiable Riemannian $\neq$ Lorentzian honesty constraint.

8 Discussion

8.1 What has and has not been shown

We have *shown* ([EST]): that four/five independent reconstruction theorems share a realization-functor form; that the finite Tannakian and finite spectral-distance cases are elementary, complete, and computable; and that the representation ladder is a well-defined mathematical object whose classical (non-condensed) sub-arrows are theorems. We have *not* shown, and explicitly do not claim: any condensed Tannakian reconstruction theorem; any recovery of a Lorentzian geometry from representation data; or that physical spacetime is, in fact, the realization of condensed data. The gap between the [EST] results and the [SPEC] hypothesis is exactly frontiers (C) and (L) of Remark 6.3, and we have tried to make that gap easy to see rather than easy to overlook.

8.2 EST/HEU/SPEC census

Claim	Status
Tannakian reconstruction $\mathcal{C} \simeq \text{Rep}_k(G_\omega)$ (Theorem 2.4)	[EST]
Finite-group reconstruction $\text{Aut}^\otimes(\omega) \cong G$ (Proposition 2.7)	[EST]
Gelfand duality; naturality (Theorem 3.1)	[EST]
Connes' spectral distance and reconstruction (Theorems 3.4 and 3.5)	[EST]
Bisognano–Wichmann: boost = modular flow (Theorem 3.7)	[EST]
Cobordism hypothesis (Theorem 3.8)	[EST]
Finite spectral distance = graph geodesic (Theorem 5.6)	[EST]
Reconstruction schema unifying the five theorems (Principle 3.9)	[HEU]
Metric read off from realization data (Heuristic 5.7)	[HEU]
Condensed Realization axiom (Definition 4.5)	[HEU]
Condensed spectral distance recovers a condensed metric (Conjecture 6.5)	[SPEC]
Spacetime is the realization of condensed data (Hypothesis 6.1)	[SPEC]
Emergent <i>Lorentzian</i> geometry from condensed data	[SPEC]

8.3 Relations and non-identifications

We reiterate the non-identifications this Part must respect:

- *Condensed Tannakian reconstruction* (the proposed generalization) is *not* classical Tannakian reconstruction (the theorem); the former is [\[SPEC\]](#), the latter [\[EST\]](#) (Definition 4.2 vs Theorem 2.4).
- *Riemannian reconstruction* (Connes) is *not* Lorentzian reconstruction (open); see Principle 6.4.
- *Bisognano–Wichmann* recovers a boost *within* a presupposed Minkowski space; it is not a from-scratch construction of spacetime.
- The condensed-mathematical approach is *not* cohesive HoTT (Schreiber) nor topos quantum theory (Isham–Döring): all three deploy sheaf/topos technology to different ends [\[23, 24\]](#), and Part V’s realization functor is external/analytic (a condensed fiber functor), not internal/synthetic.

8.4 Limitations

The principal limitations are structural, not incidental. (i) We prove nothing new about condensed categories; Definition 4.2 is a definition, not a theorem, and the reconstruction it contemplates is conjectural. (ii) All computable content is finite-dimensional; the passage to genuinely infinite (condensed) spectral triples, where the resolvent-compactness and summability conditions do real work, is exactly where the open problem lives. (iii) The Lorentzian frontier is untouched by any established result we cite; our contribution there is limited to stating the gap precisely. (iv) The composite headline claim inherits [\[SPEC\]](#) status from its weakest link and must not borrow credibility from the genuinely rigorous Riemannian/finite mathematics that surrounds it.

8.5 Outlook: hand-off to Parts VI–VII

Part VI [\[5\]](#) recomposes Parts I–V into a candidate representation-theoretic foundation for quantum gravity and must, per Principle 6.4, carry the Riemannian $\neq$ Lorentzian caveat forward without upgrade. Part VII [\[6\]](#) presents the boxed hierarchy

$$\begin{aligned} &\text{Category} \rightarrow \text{Condensed Sheaf} \rightarrow \text{Derived Object} \rightarrow \text{Representation} \\ &\rightarrow \text{Observable} \rightarrow \text{Measurement} \rightarrow \text{Emergent Spacetime} \end{aligned}$$

and inherits from this Part the precise location of the one speculative arrow (Measurement $\rightarrow$ Emergent Spacetime) and the concrete attackable target Conjecture 6.5. The most valuable next mathematical step is not physical but categorical: prove or refute Conjecture 6.5 in the static Riemannian condensed case, thereby resolving frontier (C) in isolation.

9 Conclusion

Geometry, in several of the deepest reconstruction theorems of modern mathematics, is not a primitive but a *realization*: an affine group scheme is the realization of a tensor category and a fiber functor (Tannaka); a compact space is the realization of a commutative C^* -algebra (Gelfand); a Riemannian spin manifold is the realization of a spectral triple (Connes); a Lorentz boost is the realization of a modular flow (Bisognano–Wichmann); a field theory is the realization of a fully dualizable object (cobordism hypothesis). Part V has organized these into a single representation ladder — abstract algebra $\rightarrow$ condensed object $\rightarrow$ realization $\rightarrow$ observable $\rightarrow$ geometry — and has made its established sub-arrows precise and its two computable instances (finite Tannakian reconstruction; spectral distance as a graph metric) fully explicit and code-backed.

The Part’s title claim — that *physical spacetime is the realization of condensed algebraic data* — is, honestly, a hypothesis, not a theorem. It requires two bridges that no established result supplies: a condensed Tannakian reconstruction (frontier C) and a Lorentzian upgrade of every known metric reconstruction (frontier L). We have stated the hypothesis precisely, boxed it as **[SPEC]**, isolated a well-posed Riemannian condensed target (Conjecture 6.5) that would resolve frontier (C) in isolation, and inscribed the non-negotiable Riemannian $\neq$ Lorentzian honesty constraint as a labeled principle. The value of the realization-functor viewpoint is that it converts a vague slogan (“geometry emerges”) into a definite mathematical question (“does a condensed fiber functor reconstruct a geometry, and in what signature?”) — and locates, with precision, the single arrow at which the established mathematics ends and the physics of this program begins.

10 The accompanying code

The repository directory `src/part5-geometry-as-realization-functor/` contains a self-contained Haskell package (compiling under `-Wall`) realizing the finite, decidable content of this Part:

- **Tannaka.hs** — represents a finite group by its regular representation and small irreducibles, builds the forgetful fiber functor ω , and computes the monoidal natural automorphisms $\text{Aut}^\otimes(\omega)$ directly, verifying $\text{Aut}^\otimes(\omega) \cong G$ (Proposition 2.7) for concrete groups (e.g. $\mathbb{Z}/n$, S_3).
- **Spectral.hs** — builds a finite spectral triple $(\mathcal{A}, \mathcal{H}, D)$, implements the spectral distance d_D of Definition 5.2 as the associated linear program, reproduces the two-point distance $1/|m|$ (Example 5.4), and constructs the weighted-graph triple of Construction 5.5.
- **Graph.hs** — an independent Floyd–Warshall shortest-path reference, against which **Spectral.hs**’s d_D is checked, verifying Theorem 5.6 ($d_D = d_\Gamma$) on explicit graphs.
- **Main.hs** — drives all three demonstrations and prints the reconstructed group, the two-point distance, and the graph-metric agreement.

The code is a finite *shadow* of the paper’s established content; it does not, and cannot, address frontiers (C) or (L), which are infinite-dimensional and Lorentzian respectively.

References

- [1] M. Long and The YonedaAI Collaboration, *From Manifolds to Condensed Spaces*, Part I of *Toward a Condensed Representation Theory of Physics*, YonedaAI Research Collective, 2026.
- [2] M. Long and The YonedaAI Collaboration, *Observables as Condensed Sheaves*, Part II of *Toward a Condensed Representation Theory of Physics*, YonedaAI Research Collective, 2026.
- [3] M. Long and The YonedaAI Collaboration, *Information as the Primitive Object: Condensed Cohomology and Quantum Information*, Part III of *Toward a Condensed Representation Theory of Physics*, YonedaAI Research Collective, 2026.
- [4] M. Long and The YonedaAI Collaboration, *Fields Without Background Manifolds*, Part IV of *Toward a Condensed Representation Theory of Physics*, YonedaAI Research Collective, 2026.
- [5] M. Long and The YonedaAI Collaboration, *A Representation-Theoretic Foundation for Quantum Gravity*, Part VI of *Toward a Condensed Representation Theory of Physics*, YonedaAI Research Collective, 2026 (forthcoming).
- [6] M. Long and The YonedaAI Collaboration, *Toward a Condensed Representation Theory of Physics*, Part VII (capstone) of the same series, YonedaAI Research Collective, 2026 (forthcoming).
- [7] P. Deligne and J. S. Milne, *Tannakian Categories*, in *Hodge Cycles, Motives, and Shimura Varieties*, Lecture Notes in Mathematics **900**, Springer, 1982; revised version 2022, <https://www.jmilne.org/math/xnotes/tc2022.pdf>.
- [8] N. Saavedra Rivano, *Catégories Tannakiennes*, Lecture Notes in Mathematics **265**, Springer, 1972.
- [9] P. Deligne and A. B. Goncharov, “Groupes fondamentaux motiviques de Tate mixte,” *Ann. Sci. École Norm. Sup.* **38** (2005), 1–56.
- [10] M. Levine, “Tate motives and the vanishing conjectures for algebraic K -theory,” in *Algebraic K-Theory and Algebraic Topology*, NATO ASI Series, Kluwer, 1993, 167–188.
- [11] I. M. Gelfand and M. A. Naimark, “On the imbedding of normed rings into the ring of operators in Hilbert space,” *Mat. Sbornik* **12** (1943), 197–213; see also G. K. Pedersen, *C^* -Algebras and their Automorphism Groups*, Academic Press, 1979.
- [12] A. Connes, “On the spectral characterization of manifolds,” *J. Noncommut. Geom.* **7** (2013), 1–82; arXiv:0810.2088.

- [13] A. Connes, *Noncommutative Geometry*, Academic Press, 1994; and A. Connes, “Gravity coupled with matter and the foundation of non-commutative geometry,” *Commun. Math. Phys.* **182** (1996), 155–176; arXiv:hep-th/9603053.
- [14] N. Franco and M. Eckstein, “An algebraic formulation of causality for noncommutative geometry,” *Class. Quantum Grav.* **30** (2013), 135007; arXiv:1211.3061.
- [15] M. Takesaki, *Theory of Operator Algebras I–III*, Encyclopaedia of Mathematical Sciences, Springer, 2002–2003.
- [16] J. J. Bisognano and E. H. Wichmann, “On the duality condition for a Hermitian scalar field,” *J. Math. Phys.* **16** (1975), 985–1007.
- [17] H.-J. Borchers, “On revolutionizing quantum field theory with Tomita’s modular theory,” *J. Math. Phys.* **41** (2000), 3604–3673.
- [18] J. Lurie, “On the classification of topological field theories,” in *Current Developments in Mathematics 2008*, International Press, 2009, 129–280; arXiv:0905.0465.
- [19] M. Atiyah, “Topological quantum field theories,” *Publ. Math. IHÉS* **68** (1988), 175–186.
- [20] P. Scholze, *Lectures on Condensed Mathematics*, Bonn lecture notes, 2019; stable version arXiv:2605.03658, 2026.
- [21] D. Clausen and P. Scholze, *Analytic Stacks*, lecture course, 2023 (CNRS/IHÉS); and *Lectures on Analytic Geometry* (Bonn, 2019–2020).
- [22] A. H. Chamseddine and A. Connes, “The spectral action principle,” *Commun. Math. Phys.* **186** (1997), 731–750; arXiv:hep-th/9606001.
- [23] U. Schreiber, “Differential cohomology in a cohesive ∞ -topos,” arXiv:1310.7930, 2013.
- [24] C. J. Isham and A. Döring, “Topos theory and ‘neo-realist’ quantum theory,” arXiv:0712.4003, 2007.
- [25] The YonedaAI Collaboration, *math-phy-library*: 6-module representation-stack / sheaf / HoTT / TQFT ladder (internal prior work), 2026 (representation entries, prestacks/stacks, the $\mathrm{Aut}_{[X/G]}(x) \simeq \mathrm{Stab}_G(x)$ theorem family, Realization/Equivalence axioms).
- [26] The YonedaAI Collaboration, *mathematics-physical-representation*: 5-rung motivic/coalgebraic realization ladder (internal prior work), 2026 (Betti/de Rham/Hodge/étale realizations as fiber functors; observable = per(motive)).
- [27] The YonedaAI Collaboration, *math-qg-representation-library*: 12-topic representation-theoretic dossier on quantum gravity (internal prior work), 2026 (Connes reconstruction; Bisognano–Wichmann; weakest-link composition discipline).