GrokRxiv:2026.10.yoneda-classification [math.CT] 29 Oct 2026

The Yoneda Classification Principle

Part 4 of *The Classification of Mathematical Representations of Physical Systems*

Matthew Long
The YonedaAI Collaboration, YonedaAI Research Collective
Chicago, Illinois, USA
mlong@yonedalabs.ai
https://yonedalabs.ai

October 2026

Abstract

The Yoneda embedding sends each object of a category to the functor it represents through its relationships with all other objects. This paper develops the relational characterization of mathematical representations of physical systems that the embedding supplies, and it states precisely what the embedding does not supply. The Yoneda lemma determines an object up to isomorphism by its relational profile, and the embedding is fully faithful on any locally small category of representations once the admissible morphisms are fixed. Concrete representability results for quantum supermaps show that in settings with channel-state duality the relational content is not merely formal, and the classical reconstruction theorems of Gelfand, Stone, and Doplicher and Roberts recover objects from categories of probes. Against these results stands a governing constraint: Yoneda does not, by itself, establish physical equivalence, empirical adequacy, or a classification theorem. A non-implication theorem makes the third denial precise, because full faithfulness holds in categories where isomorphism classes admit no finite-variety parametrization, so representability cannot be the source of a classification theorem. The paper separates established mathematics from conjectural physical interpretation throughout and records the open problem of unifying the supermap line with the reconstruction pattern.

1 Introduction

A mathematical representation of a physical system is a tuple of data: a state space, an observable algebra, dynamical laws, symmetry and constraint data, and composition rules. Two questions about such a tuple are easy to confuse. The first asks which equations describe the system. The second asks what relationships the system admits with every other system. Its formal home is the Yoneda embedding

$$y\colon \mathcal{C} \longrightarrow [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}], \qquad y(P) = \mathrm{Hom}_{\mathcal{C}}(-, P),$$

which sends an object P to the presheaf of its incoming morphisms. The embedding is a theorem of pure category theory: it is fully faithful on every locally small category, and an object is determined up to isomorphism by the presheaf $y(P)$.

The research program of this series asks whether mathematical representations of physical systems admit a structural classification analogous to the classification of finite simple groups. The Yoneda embedding is the natural formal home for one part of that question, because it replaces intrinsic descriptions of objects by relational data. The assignment $P \mapsto y(P)$ turns the question of

what P is into the question of which admissible relationships P has with every probe. That shift is exact mathematics. What it is not is a classification, a theory of physical equivalence, or a theory of empirical adequacy. A standing constraint of the series states the limit in one sentence, in its exact words: "Yoneda does not, by itself, establish physical equivalence, empirical adequacy, or a classification theorem." The sentence governs everything that follows.

The relational profile of a representation determines that representation up to isomorphism in a declared category of morphisms, and it determines nothing beyond that isomorphism. Full faithfulness preserves isomorphism classes without constraining how those classes are parameterized, so relational data yields no classification invariants on its own. Concrete physical recovery happens only where extra structure is assumed: channel-state duality for quantum supermaps, symmetric monoidal structure for Doplicher-Roberts reconstruction. On the positive side stand the Yoneda lemma and its corollaries, the Choi correspondence and the supermap representability it supports [1, 2], and the probe reconstruction theorems with their declared hypotheses [3, 4, 5]. On the negative side stands a non-implication theorem: full faithfulness holds in categories where isomorphism classes admit no finite-variety parametrization [6], so representability cannot be the source of a classification theorem. No theorem joins the two sides; section 9 records the open problem that would.

The classification problem of the program asks for the fundamental structural families of mathematical physics, the operations that assemble systems from them, and the equivalences that connect presentations. The Yoneda embedding lets the third question be phrased without choosing a presentation: an object is what it does to probes. A classification theorem built on relational data must supply invariants beyond representability, because representability delivers determination up to isomorphism and no parameterization of isomorphism classes. A physical reading of relational data must pass the qualification quoted above, because the morphisms defining the profiles are mathematical choices rather than physical identifications. Restricted regimes remain the available home for classification theorems.

Three distinctions organize the results. Mathematical equivalence, isomorphism of representations, is not physical equivalence, the chosen identification of systems in the world. Relational characterization establishes the first and touches the second only through the morphisms one feeds to it. Neither notion is empirical adequacy, which lives in a separate layer of the architecture.

2 Mathematical framework

The definitions below match the shared ontology of the series: an object of the category Phys is a tuple $P = (X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{K})$ comprising a state space or geometric structure X, an observable algebra or collection $\mathcal{O}$, dynamical laws $\mathcal{D}$, symmetry and constraint data $\mathcal{S}$, and composition and interaction rules $\mathcal{K}$. The series ontology writes the fifth component as $\mathcal{C}$; it is written $\mathcal{K}$ here to keep the symbol $\mathcal{C}$ available for an ambient category. The five components are a coordinate record for a representation; they need not be independent, and the schema is a candidate definition rather than a theorem.

The founding analysis of Phys names coarse-grainings, exact equivalences, embeddings, dualities, and limiting procedures as admissible relationships, and it does not fix their preservation laws or composition rules. Until those laws are fixed, Phys is a schema and not a category, and each use of the Yoneda embedding must declare which class of morphisms is in force. Accordingly, all relational statements in this paper are relative to a fixed category $\mathcal{C}$ of representations with declared morphisms. Changing the morphisms changes the relational content; nothing in the paper licenses reading a statement proved for one class of morphisms as a statement about another.

Definition 2.1 (Object schema of a representation)**.** An *object schema* is a tuple $P = (X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{K})$ where X is a state space or geometric structure, $\mathcal{O}$ is an observable algebra or collection, $\mathcal{D}$ is dynamical data, $\mathcal{S}$ is symmetry and constraint data, and $\mathcal{K}$ is composition and interaction data. A category of such tuples is *schema-faithful* when each morphism acts on the five components compatibly with the declared structure of each.

Definition 2.2 (Probe class and relational profile)**.** Let $\mathcal{C}$ be a locally small category of mathematical representations of physical systems, with the morphism classes in force declared. The *relational profile* of an object $P \in \mathrm{Ob}(\mathcal{C})$ is the presheaf

$$y(P) = \mathrm{Hom}_{\mathcal{C}}(-, P) \in [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}],$$

so that for every probe $R \in \mathrm{Ob}(\mathcal{C})$ the set $y(P)(R) = \mathrm{Hom}_{\mathcal{C}}(R, P)$ is the collection of admissible relationships from R to P, and for every morphism $f\colon R \to R'$ the action $y(P)(f)\colon \mathrm{Hom}(R', P) \to \mathrm{Hom}(R, P)$ is precomposition. For a restricted probe class $\mathcal{R} \subseteq \mathrm{Ob}(\mathcal{C})$, the *$\mathcal{R}$-restricted profile* is the family $\big(\mathrm{Hom}_{\mathcal{C}}(R, P)\big)_{R \in \mathcal{R}}$ together with the induced actions of morphisms between probes in $\mathcal{R}$.

Definition 2.3 (Relational equivalence)**.** Two objects $P, Q \in \mathrm{Ob}(\mathcal{C})$ are *relationally equivalent*, written $P \sim_{\mathrm{rel}} Q$, when $y(P) \cong y(Q)$ as presheaves on $\mathcal{C}$. They are *$\mathcal{R}$-relationally equivalent* when their restricted profiles are naturally isomorphic as functors on $\mathcal{R}$.

The two equivalence notions of the series must be kept apart. *Mathematical equivalence* is isomorphism of the tuples $P = (X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{K})$ in the ambient category of each component. *Physical equivalence* is a carefully chosen equivalence relation on representations, motivated by the observation that two mathematical descriptions may describe the same physics without being isomorphic; the choice of that relation is itself open and is load-bearing for the classification problem $\mathrm{Classify}(\mathrm{Ob}(\mathrm{Phys})/\simeq)$. Relational equivalence is a third notion, defined relative to a declared category $\mathcal{C}$, and its relationship to the first two is the subject of the next section.

Definition 2.4 (Equivalence notions and preserved layers)**.** An *identification* of two representations P and Q consists of an equivalence notion, being mathematical isomorphism or a declared physical equivalence relation, a morphism kind realizing it, and a declaration of the layers of structure the identification preserves among the symmetric monoidal, braided, monoidal, $*$-structure, algebraic, and topological layers. An identification that omits the declaration is inadmissible.

Definition 2.4 is the equivalence discipline of the series in miniature. Section 6 shows that the discipline is forced by theorems rather than chosen for tidiness: identifications that forget the monoidal or symmetric layer fail to recover groups even in the finite case.

Two properties of the framework control later claims. First, the relational profile is a function of the morphisms in force and of nothing else; a profile computed in a category of coarse-grainings is a different mathematical object from a profile computed in a category of isomorphisms. Second, the profile is relational data rather than parameter data: it assigns to each probe a set of relationships, and it does not take values in any fixed parameter space of invariants. Classification theorems in the sense of the series demand invariants whose values determine objects up to the chosen equivalence; profiles determine objects up to isomorphism by a different mechanism entirely.

Remark 2.5 (Finite-dimensional quantum systems). Fix $n \geq 1$ and let $\mathcal{C}_n$ be the category whose objects are pairs (H, A) of a finite-dimensional complex Hilbert space H with $\dim H \leq n$ and a unital $*$-subalgebra $A \subseteq B(H)$, with morphisms the $*$-homomorphisms. The tuple data of the series is recovered by taking $X = H$ as state space, $\mathcal{O} = A$ as observables, $\mathcal{D}$ the specified Hamiltonian or semigroup dynamics when present, $\mathcal{S}$ the specified symmetry action when present, and $\mathcal{K}$ the

tensor product of systems. The relational profile of (H, A) records, for each probe (H', A'), the set of unital $*$-homomorphisms $A' \to A$ together with their compatibilities. This category is locally small, and the Yoneda embedding applies to it as to any other.

Remark 2.6 (Classical probe systems). Let $\mathcal{C}_{\mathrm{cl}}$ be the category of compact Hausdorff spaces with continuous maps, regarded as the geometric side of classical representations, and let the observable side be the commutative C^*-algebra $C(X)$ of continuous complex-valued functions. A probe R contributes the set $\mathrm{Hom}(R, X)$ of continuous maps, equivalently the unital $*$-homomorphisms $C(X) \to C(R)$ through Gelfand duality. Section 5 shows that on the commutative side the probe data reconstructs the object itself.

Remark 2.7 (Process categories). Fix finite-dimensional systems and let $\mathcal{C}_{\mathrm{ch}}$ have as objects the completely positive trace-preserving maps $\Phi\colon M_n \to M_m$ for varying n and m. A morphism from $\Phi\colon A \to B$ to $\Psi\colon C \to D$ is a pair of channels $(E\colon C \to A,\ L\colon B \to D)$ with $\Psi = L \circ \Phi \circ E$; composition is componentwise, and the identity on Φ is $(\mathrm{id}, \mathrm{id})$. These pairs are the pre- and post-processings without auxiliary memory. General supermaps admit ancillary systems that bridge input and output, and the representability theorem of Section 4 covers that wider setting. Supermaps are the transformations of the processes this category treats as objects, and the relational content in this setting is concretely computable through channel-state duality, because a channel is a state of the composed system.

3 Relational characterization

The Yoneda lemma and its corollaries [7] hold in every locally small category, and their application to Phys is conditional on the declaration of morphisms recorded in section 2.

Theorem 3.1 (Yoneda lemma)**.** *Let $\mathcal{C}$ be a locally small category and let $A \in \mathrm{Ob}(\mathcal{C})$. For every functor $F\colon \mathcal{C}^{\mathrm{op}} \to \mathbf{Set}$ the evaluation map*

$$\mathrm{ev}_A\colon \mathrm{Nat}\left(\mathrm{Hom}_{\mathcal{C}}(-, A), F\right) \longrightarrow F(A), \qquad \alpha \longmapsto \alpha_A(\mathrm{id}_A),$$

is a bijection.

Proof. Injectivity first. Suppose $\alpha, \beta\colon \mathrm{Hom}(-, A) \to F$ satisfy $\alpha_A(\mathrm{id}_A) = \beta_A(\mathrm{id}_A)$. For any object X and any $f \in \mathrm{Hom}(X, A)$ the naturality square at $f\colon X \to A$ gives $\alpha_X(f) = F(f)(\alpha_A(\mathrm{id}_A))$ and likewise for β, so $\alpha_X = \beta_X$ for every X, and $\alpha = \beta$.

Surjectivity next. Given $a \in F(A)$, define $\alpha^a_X\colon \mathrm{Hom}(X, A) \to F(X)$ by $\alpha^a_X(f) = F(f)(a)$. Naturality holds: for $g\colon Y \to X$ and $f \in \mathrm{Hom}(X, A)$ one has $\alpha^a_Y(f \circ g) = F(f \circ g)(a) = F(g)(F(f)(a)) = F(g)(\alpha^a_X(f))$, which is exactly the required square. Finally $\alpha^a_A(\mathrm{id}_A) = F(\mathrm{id}_A)(a) = a$, so $\mathrm{ev}_A(\alpha^a) = a$. The two constructions are mutually inverse, and local smallness guarantees that each $\mathrm{Hom}(X, A)$ is a set on which F acts. □

Corollary 3.2 (Full faithfulness)**.** *For all objects A, B of a locally small category $\mathcal{C}$ the map*

$$\mathrm{Hom}_{\mathcal{C}}(A, B) \longrightarrow \mathrm{Nat}\left(y(A), y(B)\right), \qquad \varphi \longmapsto \left(g \mapsto \varphi \circ g\right),$$

is a bijection. In particular y is faithful and full, and $y(A) \cong y(B)$ holds if and only if $A \cong B$.

Proof. Apply Theorem 3.1 with $F = \mathrm{Hom}(-, B)$; the bijection sends the natural transformation α to $\alpha_A(\mathrm{id}_A)$, and its inverse sends φ to $(g \mapsto \varphi \circ g)$. Faithfulness and fullness are the statement that this map is bijective on hom-sets. If $\theta\colon y(A) \to y(B)$ is a natural isomorphism, the corresponding $\varphi = \theta_A(\mathrm{id}_A)$ is an isomorphism in $\mathcal{C}$, with inverse induced by θ^{-1}; conversely an isomorphism $A \cong B$ induces one of profiles by precomposition. □

Corollary 3.2 is the exact content of the relational shift. An object P is fully described, up to isomorphism in $\mathcal{C}$, by the totality of its admissible relationships $\mathrm{Hom}(R, P)$ with all probes R, together with the compatibility of those relationships along morphisms of probes. No intrinsic presentation of P is needed to recover P up to isomorphism.

Proposition 3.3 (Probe content of a profile)**.** *Let $P \in \mathrm{Ob}(\mathcal{C})$ with $\mathcal{C}$ locally small. For each probe R the value $y(P)(R) = \mathrm{Hom}(R, P)$ is the set of admissible relationships from R to P, and the action of $y(P)$ along $f\colon R \to R'$ is the composition law of $\mathcal{C}$. Consequently the profile $y(P)$ is determined by, and determines, the system's interaction with every probe in $\mathcal{C}$, and two systems have isomorphic profiles exactly when those interactions coincide up to relabelling by isomorphisms of the probes' targets.*

Proof. The value statement is the definition of y. The action statement is precomposition, which is the composition law of $\mathcal{C}$ by definition. The last claim is Corollary 3.2. □

Corollary 3.4 (Relational equivalence is isomorphism)**.** *For a declared category $\mathcal{C}$ of representations one has $P \sim_{\mathrm{rel}} Q$ if and only if P and Q are mathematically equivalent in the sense of isomorphism in $\mathcal{C}$. Relational equivalence therefore refines or coincides with mathematical equivalence according to the morphisms in force, and it implies physical equivalence only when the declared morphisms have themselves been justified as physical equivalences.*

Proof. The first sentence is Corollary 3.2. The second holds because $\sim_{\mathrm{rel}}$ is by definition isomorphism in the declared category, while physical equivalence is a chosen relation on representations whose correctness is a separate question; the two coincide only if the morphisms of $\mathcal{C}$ were chosen to encode exactly the justified physical identifications. □

Relational characterization is rigorous, and it delivers mathematical equivalence in the category one declares. It reaches physical equivalence only through a justified choice of morphisms, and no such choice is fixed by the embedding itself. The denial concerning classification takes the form of a non-implication: full faithfulness of the embedding does not imply that isomorphism classes admit separating invariants.

The presheaf acts on each probe morphism $f\colon R \to R'$ by precomposition, sending a relationship in $\mathrm{Hom}(R', P)$ to its composite with f in $\mathrm{Hom}(R, P)$, so relational data is only as fine as the probe class allows. Two representations whose hom-sets agree on a subcategory $\mathcal{R} \subseteq \mathcal{C}$ are indistinguishable to those probes, and restricting the probe class can destroy separation entirely, as Proposition 3.6 records. Distinguishability through profiles is therefore bounded by what the chosen probes can ask, and the fragment of probes is part of every determination claim.

3.1 Restricted probes

The full profile uses every object of $\mathcal{C}$ as a probe. In practice one computes with a probe class $\mathcal{R}$, and the question is what the restriction loses.

Remark 3.5 (A small computation). Representable presheaves distinguish non-isomorphic objects on finite categories. Let $\mathcal{C}$ have objects A, B and morphisms $\mathrm{id}_A, \mathrm{id}_B$, and one non-identity morphism $f\colon A \to B$. The presheaves evaluate on objects to $y(A)(A) = \{\mathrm{id}_A\}$, $y(A)(B) = \varnothing$, $y(B)(A) = \{f\}$, and $y(B)(B) = \{\mathrm{id}_B\}$. By Corollary 3.2 the arrow $f \in \mathrm{Hom}(A, B)$ induces a unique natural transformation $\theta\colon y(A) \to y(B)$ with components $\theta_A(\mathrm{id}_A) = f$ and $\theta_B\colon \varnothing \to \{\mathrm{id}_B\}$. Naturality on the probe morphism $f\colon A \to B$ demands $\theta_A \circ y(A)(f) = y(B)(f) \circ \theta_B$; both composite maps go from $y(A)(B) = \mathrm{Hom}(B, A) = \varnothing$ to $\mathrm{Hom}(A, B) = \{f\}$ and therefore coincide vacuously. The values of the two profiles on the probe B differ, being $\varnothing$ and $\{\mathrm{id}_B\}$, so $y(A) \ncong y(B)$, matching $A \ncong B$.

Proposition 3.6 (What restriction loses)**.** *Let $\mathcal{R} \subseteq \mathrm{Ob}(\mathcal{C})$ be a probe class. Isomorphic restricted profiles do not imply isomorphic objects in general. If the restricted Yoneda comparison $y_{\mathcal{R}} \colon \mathcal{C} \to [\mathcal{R}^{\mathrm{op}}, \mathbf{Set}]$ is fully faithful, then $\mathcal{R}$-relationally equivalent objects are isomorphic, and the argument of Corollary 3.2 applies verbatim.*

Proof. For the first claim take $\mathcal{C}$ with three objects A, B, C, no morphisms between distinct objects, and $\mathcal{R} = \{C\}$. Then $\mathrm{Hom}(C, A) = \mathrm{Hom}(C, B) = \varnothing$, so the restricted profiles of A and B agree, while $A \ncong B$. For the second claim, full faithfulness of $y_{\mathcal{R}}$ means $\mathrm{Hom}(A, B) \cong \mathrm{Nat}(y_{\mathcal{R}}(A), y_{\mathcal{R}}(B))$ for all A, B, so a natural isomorphism of restricted profiles induces an isomorphism of objects exactly as in Corollary 3.2. □

Proposition 3.6 makes the probe class part of the hypothesis. A profile computed on an arbitrary small probe class carries no more content than that class exposes, and the condition that the comparison be fully faithful is the standard density requirement on $\mathcal{R}$. Any physical claim about determination by probes must therefore state the probes.

3.2 Composition of systems

Physical systems compose, through tensor products, direct sums, and the other constructions that generate complex systems from simple ones. The behavior of relational profiles under composition is therefore a question the classification program cannot avoid.

Proposition 3.7 (Profiles respect limits)**.** *Let $\mathcal{C}$ be a locally small category in which a diagram D has a limit $L = \lim D$. Then $y(L)$ is the limit of the diagram $y \circ D$ in $[\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$. In particular y preserves terminal objects and binary products: $y(A \times B) \cong y(A) \times y(B)$ canonically whenever the product exists.*

Proof. Limits of presheaves are computed pointwise, so for each probe R one must show

$$\mathrm{Hom}(R, \lim D) \cong \lim \mathrm{Hom}(R, D),$$

which is the defining universal property of the limit applied to the probe R. For binary products the statement reads $\mathrm{Hom}(R, A \times B) \cong \mathrm{Hom}(R, A) \times \mathrm{Hom}(R, B)$, which is the universal property of the product. □

Tensor products behave differently. There is in general no canonical map relating $\mathrm{Hom}(R, P \otimes Q)$ to $\mathrm{Hom}(R, P) \otimes \mathrm{Hom}(R, Q)$, and the profile of a composite system is not assembled from the profiles of its factors by any universal construction available to the embedding. Determination of composites from relational data is therefore not a formal consequence of Yoneda; where it holds, it holds by a further theorem. The factorisation results of section 4 are exactly such a theorem for the composition of processes under a causal-order hypothesis, and the decomposition theory of representation categories is such a theory for tensor structure. This asymmetry is a design fact of the relational program: limits are transparent to probes, monoidal composition is not.

3.3 The classification question

Definition 3.8 (Classification by invariants)**.** Let $\mathcal{C}$ be a category of representations and let S be a parameter set. An *invariant* is a map $I \colon \mathrm{Ob}(\mathcal{C}) \to S$ constant on isomorphism classes. The invariant I *classifies* $\mathcal{C}$ up to isomorphism when $I(P) = I(Q)$ implies $P \cong Q$; a family of invariants classifies when the joint map does. A classification is *of finite-variety type* over a field k when the parameter set S is the set of k-points of a finite-dimensional algebraic variety over k. A class of

objects is *parametrized by a finite-dimensional variety over k* when such a classifying invariant of finite-variety type exists.

In the series vocabulary a classifying invariant is an invariant of this kind whose value determines the object up to the chosen equivalence, in the manner of a Dynkin diagram for semisimple Lie algebras or a dimension vector for finite-dimensional C^*-algebras. The Yoneda embedding supplies no such data: $y(P)$ is a functor rather than a value in a parameter space, and deciding $y(P) \cong y(Q)$ is in general as hard as deciding $P \cong Q$. The temptation the governing constraint guards against is to read determination by relationships as classification. The next theorem separates the two notions formally; its witness is the boundary of representability itself rather than a counterexample to any published claim.

Theorem 3.9 (Representability does not imply classifiability)**.** *Let $n \geq 3$, let $S = k[x_{ij}]$ be the polynomial ring over a field k in the entries of a generic $n \times n$ matrix, let $R = S/(\det(x_{ij}))$ be the coordinate ring of the generic determinant hypersurface, and let $\mathcal{C}$ be the category of finitely generated maximal Cohen-Macaulay R-modules of rank two, with R-module homomorphisms as morphisms. The objects of $\mathcal{C}$ admit no parametrization by a finite-dimensional variety over k in the sense of Definition 3.8, while y is fully faithful on $\mathcal{C}$. Full faithfulness of the Yoneda embedding therefore does not imply the existence of a classification theorem.*

Proof. The category $\mathcal{C}$ is locally small, so Theorem 3.1 and Corollary 3.2 apply: y is fully faithful on $\mathcal{C}$. The non-parametrization is the classification result of Buchweitz and Leuschke [6] for modules over the generic determinant of an $n \times n$ matrix with $n \geq 3$; the bound on n is part of their setting, since the small cases are tame. Their conclusion is exact: the Cohen-Macaulay representation theory of the generic determinant is wild in the tame-wild dichotomy, and in their words it is quite wild, since even in rank two the isomorphism classes cannot be parametrized by a finite-dimensional variety over the coefficients. Rank two is precisely the case their statement covers, which is why $\mathcal{C}$ is defined there. Both statements hold in the same category: a classification in the sense of Definition 3.8 is absent for $\mathcal{C}$, while full faithfulness of y on $\mathcal{C}$ is complete. If full faithfulness implied classifiability, no such category would exist; the category exists; so the implication fails. □

Remark 3.10 (Scope of the non-implication)*.* Theorem 3.9 denies an implication; it does not deny that relational data is useful. In tame regimes relational methods convert into invariants, and Sections 4 and 5 exhibit settings where they do. The theorem says that the conversion is a further theorem, with hypotheses, and that it cannot be smuggled in as a consequence of representability. The wider wildness literature reinforces the point: every finite-dimensional algebra over an algebraically closed field is derived tame or derived wild, and wildness is preserved under degeneration [8], so the failure pattern is structural rather than isolated.

The three denials of the governing constraint now have their mathematical content. Physical equivalence is a choice of morphisms or relations that the embedding does not make. Empirical adequacy is a property of how a representation meets observation, a separate layer of the architecture from mathematical classification. A classification theorem is a supply of separating invariants, and Theorem 3.9 shows that representability is logically independent of such a supply. What Yoneda does establish is Corollary 3.4: objects are determined by their relationships, up to the isomorphism encoded in the declared morphisms.

3.4 Determination and classification

Structural knowledge of a family falls into one of four regimes, recorded in Table 1; the Yoneda embedding belongs to the third.

Table 1: Four regimes of structural knowledge. The Yoneda embedding supplies determination; classification content comes from the invariants of each family; the wild regimes forbid finite-variety classification outright.

Regime	Example	Invariant	What is missing
Complete and parameterized	Semisimple Lie algebras	Dynkin diagram	Nothing in scope
Complete without parameters	Finite simple groups	Atlas data	A uniform parameterization
Determination only	Any locally small category	Relational profile	Separating invariants
Forbidden in the stated generality	Modules over the generic determinant	None of finite-variety type	Classification itself

The second regime is the reminder that even a completed classification can lack a uniform parameterization, which is why the classification problem of this series asks for families and constructions rather than for a single list. The third regime is where this paper lives: determination is available everywhere, and classification content is a further achievement of particular families. The fourth regime is the boundary drawn by Theorem 3.9: in wild settings determination persists and classification does not, and no amount of representability changes that. A relational program for classification must therefore work through the second regime, converting determination into family-specific invariants wherever tame structure allows.

4 Supermap representability

Formal category theory gives relational profiles no matter what the objects are. The question of whether relational content is ever concretely computational, rather than formal, has a positive answer in one corner of quantum theory: quantum supermaps, the transformations of quantum channels. The representability theorems of 2024 to 2026 show that in settings with channel-state duality, categorical supermaps are concretely representable, and that Yoneda representability underlies factorisation theorems for them [1, 2].

A *supermap* is a transformation Ψ that assigns to each channel $\Phi\colon M_n \to M_m$ of an admissible class a channel $\Psi(\Phi)\colon M_{n'} \to M_{m'}$, compatibly with composition of channels and with the tensor product of systems. Supermaps are therefore second-order processes: processes whose inputs and outputs are themselves processes. The relational question for them is concrete. Which transformations count as processes in their own right, and can the answer be read off from the relationships the channel admits? The channel-state duality answers both questions at once, because it identifies a channel with a state of the composed system.

Let M_n denote the $n \times n$ complex matrices, and let $\Phi\colon M_n \to M_m$ be a linear map. The *Choi matrix* of Φ [9] is

$$J(\Phi) \;=\; \sum_{i,j=1}^{n} E_{ij} \otimes \Phi(E_{ij}) \;\in\; M_n \otimes M_m,$$

where E_{ij} are the matrix units of M_n. A linear map is completely positive if and only if $J(\Phi) \geq 0$, and it is trace preserving if and only if $\mathrm{Tr}_m\, J(\Phi) = I_n$, the partial trace over the second tensor factor. The normalized Choi matrix of a channel is a state on $M_n \otimes M_m$, and this correspondence between channels and states is *channel-state duality*.

Proposition 4.1 (Choi correspondence and composition)**.** *For fixed $n, m \geq 1$ the assignment $\Phi \mapsto J(\Phi)$ is a linear bijection $L(M_n, M_m) \to M_n \otimes M_m$. It restricts to a bijection between completely positive maps and positive semidefinite Choi matrices, and Φ is trace preserving if and only if $\mathrm{Tr}_m J(\Phi) = I_n$. If $\Phi\colon M_n \to M_m$ and $\Psi\colon M_m \to M_p$ are linear, then*

$$J(\Psi \circ \Phi)_{(i\alpha),(j\beta)} \;=\; \sum_{k,l=1}^{m} J(\Phi)_{(ik),(jl)}\, J(\Psi)_{(k\alpha),(l\beta)}, \tag{1}$$

where composite indices label rows and columns of $M_n \otimes M_m$ and $M_m \otimes M_p$. Composition of channels is therefore a contraction of Choi indices along the shared system. In the link-product notation of quantum information the same identity appears as $J(\Psi \circ \Phi) = J(\Phi) * J(\Psi)$, where the standard convention places a partial transpose on the contracted factor; the component formula displayed here pairs row with row and column with column indices of the shared system, which absorbs that transposition into the index placement.

Proof. The assignment is linear in Φ by construction. Write $\Phi^{ij,kl}$ for the elementary linear maps $\Phi^{ij,kl}(X) = \mathrm{Tr}(E_{ji}X)\,E_{kl}$; these form a basis of $L(M_n, M_m)$, and $J(\Phi^{ij,kl}) = E_{ij} \otimes E_{kl}$. The elements $E_{ij} \otimes E_{kl}$ form a basis of $M_n \otimes M_m$, so the assignment carries a basis to a basis between vector spaces of the same dimension n^2m^2 and is therefore bijective. Complete positivity corresponds to positivity of $J(\Phi)$, which is the standard matrix-unit characterization. For the trace condition, the partial trace $\mathrm{Tr}_m J(\Phi)$ has entries $\mathrm{Tr}(\Phi(E_{ij}))$, and Φ is trace preserving exactly when $\mathrm{Tr}(\Phi(E_{ij})) = \mathrm{Tr}(E_{ij}) = \delta_{ij}$ for all i, j, which says $\mathrm{Tr}_m J(\Phi) = \sum_i E_{ii} = I_n$. For the composition identity, write $\Phi(E_{ij}) = \sum_{k,l} c_{kl}^{ij} E_{kl}$. Then $J(\Phi)_{(ik),(jl)} = c_{kl}^{ij}$ and $J(\Psi)_{(k\alpha),(l\beta)} = \Psi(E_{kl})_{\alpha\beta}$. The right-hand side of the claimed identity is $\sum_{k,l} c_{kl}^{ij} \Psi(E_{kl})_{\alpha\beta} = \Psi(\Phi(E_{ij}))_{\alpha\beta}$, which is $J(\Psi \circ \Phi)_{(i\alpha),(j\beta)}$ by definition. □

Remark 4.2 (Worked instance of the contraction). Take $n = m = 2$, let Φ be the completely depolarizing channel $\Phi(X) = \mathrm{Tr}(X)\, I/2$, and let Ψ be the identity channel. Then $J(\Phi) = \frac{1}{2} I \otimes I$ and $J(\Psi) = |\Omega\rangle\langle\Omega|$ with $|\Omega\rangle = \sum_k |kk\rangle$ the unnormalized maximally entangled vector; both are positive semidefinite with $\mathrm{Tr}_m J = I$. The contraction of (1) reads $\sum_{k,l} J(\Phi)_{(ik),(jl)} \delta_{k\alpha}\delta_{l\beta} = \frac{1}{2}\delta_{ij}\delta_{\alpha\beta}$, which is $J(\Psi \circ \Phi) = J(\Phi)$, since $\Psi \circ \Phi = \Phi$. The experiment of section 7 checks the same identity on sampled channels with random Kraus operators.

Proposition 4.1 is the concrete setting in which channel-state duality is exact: the channel is a state on the composed system, and composition is an index contraction. The category of Remark 2.7 records the fragment of process morphisms without auxiliary memory; the theorem below covers supermaps in the full sense of [1, 2], where ancillary systems may bridge input and output, and the presheaf of the restricted category is the fragment of that structure on pre- and post-processings.

Theorem 4.3 (Supermap representability)**.** *Let $\mathcal{T}$ be a theory of processes whose processes form a symmetric monoidal category under sequential and parallel composition, and suppose $\mathcal{T}$ is equipped with a channel-state duality: a natural identification of each process $\Phi\colon A \to B$ with a state of the composite system $A \otimes B$, preserved by composition. Then the categorical supermaps of $\mathcal{T}$ are concretely representable through that duality: every supermap corresponds to a process of the composite system, and representability yields the factorisation theorems for causally ordered supermaps.*

Proof. The statement is the representability lemma established in [1]; the factorisation theorems are proved in [2]. We record the result as a registered input of the series and do not reproduce the arguments there. Proposition 4.1 verifies the hypothesis in the finite-dimensional matrix setting:

the Choi matrix identifies each channel with a state of the composite system, and composition of channels matches contraction of Choi indices, so the duality is preserved by composition as the hypothesis requires. □

Remark 4.4 (What the theorem shows). Theorem 4.3 is the one recorded setting where relational representability is concretely computational: the relationships between processes are themselves processes, and factorisation statements follow. The hypotheses travel with the theorem. Channel-state duality of the required kind is a property of particular frameworks, and no claim is made here that every setting of interest has it, or that the result extends from quantum processes to mathematical representations of physical systems at large. Wilson, Hefford, and Hoffreumon [1] apply the same lemma to generalised circuit theories, including boxworld and real quantum theory, which shows the hypothesis is checkable theory by theory rather than automatic.

The supermap setting supports one precise reading of the relational program: once the morphisms are fixed as process transformations and a channel-state duality is available, the profile of a process is not formal data but computable data. That is genuine content, and it is narrower than the slogan of a Yoneda lemma for physics. Section 9 records how much narrower.

4.1 Factorisation theorems

The representability statement is not the end of the supermap line. The same authors use it to prove factorisation theorems [2], and the shape of those theorems is the clearest existing instance of relational structure producing concrete process structure.

The 2024 work [2] identifies morphisms of strong profunctors as a categorification of quantum supermaps, which places these black-box generalisations of diagrams with holes inside the theory of profunctor optics. From this vantage it constructs abstract logical connectives for supermaps, including tensor products and negations, in a theory-independent setting. Three structural features of quantum supermaps are then modelled by these connectives alone: black-box indefinite causal order, black-box definite causal order, and the factorisation of definitely causally ordered supermaps into concrete circuit diagrams. The 2026 work [1] states the mechanism as the lemma of Theorem 4.3, and the lemma eliminates guesswork in defining supermaps for a given theory.

Two features of this line matter for the classification program. The factorisation theorems are classification-like results for a restricted class: under a causal-order hypothesis, supermaps decompose into circuit data, which is determination by constituents in the sense of decomposition rather than determination by probes. The hypotheses are physical and checkable per theory, and [1] applies them to generalised theories, including boxworld and real quantum theory, rather than to quantum theory alone. The relational principle has so far earned its keep as a method with hypotheses, not as a universal law.

5 Reconstruction from probes

A second family of results shares the shape of the relational program: reconstruction theorems, which recover an object from a category of probes attached to it. The pattern is older than the supermap line and its settings are exact.

Theorem 5.1 (Gelfand duality)**.** *Let* **KHaus** *be the category of compact Hausdorff spaces with continuous maps, and let* **CCStar** *be the category of unital commutative complex C^*-algebras with unital $*$-homomorphisms. The assignment $X \mapsto C(X)$ and $A \mapsto \mathrm{Sp}(A)$ extends to an equivalence of categories* $\mathbf{KHaus}^{\mathrm{op}} \simeq \mathbf{CCStar}$*, a contravariant equivalence between* **KHaus** *and* **CCStar***;*

equivalently, every unital commutative C^*-algebra is isometrically $*$-isomorphic to $C(X)$ for X its character space. More generally, the Gelfand-Naimark theorem [3] embeds every C^*-algebra in the bounded operators of a Hilbert space, and the commutative case is the fully classified one.

Proof. The proofs are the standard ones [3]: characters of A separate points and give the Gelfand transform $A \to C(\mathrm{Sp}(A))$, which is an isometric $*$-isomorphism for unital commutative A; the inverse sends X to $C(X)$ with $\mathrm{Sp}(C(X)) \cong X$ by evaluation. The general embedding result is the Gelfand-Naimark theorem. □

Theorem 5.2 (Stone duality)**.** *Let* **Bool** *be the category of Boolean algebras with homomorphisms, and let* $\mathbf{KHaus}_0$ *be the category of zero-dimensional compact Hausdorff spaces with continuous maps. The Stone space construction $B \mapsto \mathrm{St}(B)$ and the algebra of clopen sets $X \mapsto \mathrm{Clop}(X)$ extend to an equivalence of categories $\mathbf{Bool}^{\mathrm{op}} \simeq \mathbf{KHaus}_0$, a contravariant equivalence between* **Bool** *and* $\mathbf{KHaus}_0$ *[4].*

Proof. Ultrafilters of a Boolean algebra B form a compact zero-dimensional space with the topology generated by the sets of ultrafilters containing each element; the clopen algebra of that space recovers B. The two constructions are mutually inverse up to natural isomorphism [4]. □

Theorem 5.3 (Doplicher-Roberts reconstruction)**.** *Let $\mathcal{T}$ be a symmetric tensor C^*-category with conjugates, direct sums, subobjects, and* $\mathrm{End}(\mathbf{1}) \cong \mathbb{C}$*. Then $\mathcal{T}$ is equivalent, as a symmetric tensor $*$-category, to the category of finite-dimensional unitary representations of a compact supergroup G, and G is unique up to isomorphism [5]. The hypotheses are not decorative: symmetric structure, conjugates, direct sums, subobjects, and the scalar endomorphism condition are all used.*

Proof. The reconstruction is that of [5], applied in its original setting to the representations selected by the DHR criteria of a local net [10]; the uniqueness statement is part of the same theorem. The hypotheses of the statement are exactly those listed; dropping any of them invalidates the conclusion, and section 6 records the constraint that dropping the symmetric layer is provably insufficient in any case. □

Remark 5.4 (The DHR setting)*.* The Doplicher-Roberts theorem arises from a net of local observables on spacetime: representations of the net are selected by their behavior on the causal complements of regions [10], and the selected representations form exactly the symmetric tensor C^*-category of Theorem 5.3. The probes are the localized endomorphisms, and the recovered object is the compact supergroup implementing the superselection sectors. When the parity twist is trivial, the recovered object is an ordinary compact group; the supergroup form records a non-trivial parity in the symmetric structure. The finite-group statements of section 6 are the trivial-twist case, where a symmetric representation category recovers the group itself. This is the sense in which a physical setting supplies a probe category by construction; it is also the reason the hypotheses are physical, not just technical.

Each theorem has the same shape: a category of probes determines the object. For Gelfand, the probes are characters; for Stone, the probes are ultrafilters; for Doplicher-Roberts, the probes are the localized endomorphisms of a net of observables. In each case the recovering functor is contravariant, and in each case the statement is an equivalence of categories, which is stronger than determination of a single object. Table 2 collects the data.

Proposition 5.5 (Probe reconstruction on the commutative side)**.** *Let A be a unital commutative C^*-algebra with character space $X = \mathrm{Sp}(A)$. The Gelfand transform $a \mapsto \hat{a}$ is the map that assembles the probe values of a, that is, the numbers $\chi(a)$ for the characters $\chi \in \mathrm{Hom}(A, \mathbb{C})$, and the statement*

Table 2: Probe reconstruction islands: each row recovers an object from a category of probes, with the hypotheses that make the recovery exact.

Result	Probe category	Recovered object	Hypotheses in force
Gelfand duality	Characters of a commutative algebra	Compact Hausdorff spectrum	Unital commutative C^*-algebra
Stone duality	Ultrafilters of a Boolean algebra	Zero-dimensional compact space	Boolean algebra
Doplicher-Roberts	Localized endomorphisms of a net	Compact supergroup	Symmetric tensor C^*-category with conjugates

of Theorem 5.1 is that these probe values determine A up to isometric $*$-isomorphism. The probe data is thus the relational data of A against the one-dimensional representations, and the reconstruction theorem upgrades the determination to an equivalence of categories.

Proof. By definition $\hat{a}(\chi) = \chi(a)$, so the Gelfand transform is the assembly of probe values. Theorem 5.1 states that the transform is an isometric $*$-isomorphism onto $C(X)$, which is the determination claim; the contravariant equivalence of categories is the upgrade. □

The relationship of these theorems to Theorem 4.3 is the open question of section 9. Structurally, each is a relational determination result, and Proposition 5.5 shows how literally the probe reading can be taken in the commutative case. Whether the supermap representability line and these reconstruction theorems are instances of one principle is not settled by the present evidence. The resemblance is a research observation; no unifying theorem is claimed in this paper, and the qualification of the governing constraint applies to every attempt to synthesize them.

The boundaries of the islands are as informative as their statements. Gelfand duality classifies the commutative unital part of the operator-algebra family and stops at noncommutativity, where the Gelfand-Naimark theorem supplies an embedding but no classification. Stone duality reaches exactly the zero-dimensional compact spaces and the Boolean algebras, and no further. Doplicher-Roberts reconstruction is available only under symmetric tensor hypotheses, and section 6 shows that weakening the symmetric layer destroys the recovery of the group in general. Each theorem is an equivalence of categories with a domain of validity, and the domains do not presently overlap into one principle. A further distinction runs through all three: they determine objects by an equivalent category of probes, which is not parametric classification in the sense of Definition 3.8, since the probe objects are as varied as the objects determined.

Remark 5.6 (Finite Stone duality). A finite instance makes the probe reading concrete. Let X be a finite discrete space of n points. Its clopen algebra is the full power set 2^X, a Boolean algebra with 2^n elements, and the Stone space of 2^X is the set of ultrafilters, one per point of X, recovering X exactly. The probes are the ultrafilters, equivalently the two-valued homomorphisms $2^X \to 2$, and the recovery is the statement that a finite set is determined by its maps into the two-element set. This is the smallest instance of Theorem 5.2, and the experiment of section 7 computes its probe data directly.

6 Equivalence and preserved structure

Relational data distinguishes objects up to isomorphism in the declared category. Which layers of structure that isomorphism preserves is not a bookkeeping question: there are theorems forbidding

the naive answers.

Theorem 6.1 (Etingof-Gelaki constraint)**.** *There exist non-isomorphic finite groups G and H whose categories* $\mathrm{Rep}(G)$ *and* $\mathrm{Rep}(H)$ *of finite-dimensional complex representations are equivalent as monoidal categories, with the commutativity constraint ignored, while the symmetric structures of the two categories differ. A monoidal equivalence of representation categories therefore does not recover the group; the symmetric monoidal layer is required, and any identification that relies on the monoidal layer alone is provably insufficient.*

Proof. The existence of such isocategorical groups is proved in [11]: the paper constructs non-isomorphic groups whose representation categories over $\mathbb{C}$ are equivalent as tensor categories without regard for the commutativity constraint, and records at the same time that symmetric tensor categories of representations do determine a finite group. The negative content for this series is that the monoidal layer is not enough. □

Theorem 6.1 is a constraint on every equivalence proposed in this series: any equivalence used to identify systems must state which layer of structure it preserves. The constraint reaches relational profiles directly.

The construction behind the theorem explains where the failure lives. Starting from a finite group G and a suitable cohomological twist of the symmetry, one obtains non-isomorphic groups whose representation categories carry the same monoidal structure while the symmetric structure differs. The group data is therefore not readable from the monoidal layer even when that layer is read in full, and no choice of relational probes computed in a category of monoidal equivalences can recover it. This is the precise reason the declaration of preserved layers is mandatory rather than advisory.

Proposition 6.2 (Preserved-layer discipline for profiles)**.** *Let $\mathcal{C}$ be a category of representations whose morphisms preserve a layer L of structure and forget a layer L'. If $P \sim_{\mathrm{rel}} Q$ in $\mathcal{C}$, then P and Q are equivalent at layer L, and no claim about layer L' follows. In particular, profiles computed in a category of monoidal equivalences that ignore the commutativity constraint cannot support identifications that require the symmetric layer, and every use of the relational profile must declare the layers its morphisms preserve.*

Proof. By Corollary 3.4 relational equivalence is isomorphism in $\mathcal{C}$, which preserves exactly the structure the morphisms of $\mathcal{C}$ preserve. Theorem 6.1 exhibits non-isomorphic groups whose representation categories are equivalent as monoidal categories, so no inference about the forgotten symmetric layer is available from layer-L equivalence alone. The declaration requirement is the same constraint applied to the morphisms of $\mathcal{C}$. □

Corollary 6.3 (Constraint inheritance)**.** *The Yoneda embedding does not weaken the structure-preserving requirement on equivalences. Every relational identification inherits, through its morphisms, the preserved-layer declaration of the equivalence notion in force; a profile computed for one declaration is not evidence for an identification requiring another.*

Proof. Proposition 6.2 with L and L' read as the declared and forgotten layers. □

A computable shadow sits two decategorification steps below Theorem 6.1: the fusion ring of a representation category. Let $\mathrm{Rep}(G)$ have irreducible objects $V_1, \dots, V_r$ with fusion multiplicities N_{ij}^k defined by $V_i \otimes V_j \cong \bigoplus_k N_{ij}^k V_k$. The resulting fusion ring records only the tensor multiplicities, not the associator of the monoidal category.

Table 3: Layers of structure an identification may preserve, and what each layer recovers. The second row is the content of the Etingof-Gelaki constraint: the monoidal layer loses the group.

Layer	Recovers	Status
Symmetric monoidal	The finite group of a representation category	Sufficient for finite groups
Monoidal (tensor)	Tensor structure without the commutativity constraint	Provably insufficient for the group
Fusion ring	Fusion multiplicities and dimensions only	Insufficient (Proposition 6.4)
Algebraic and $*$-structure	The operator-algebraic presentation	Independent of the tensor layers

Table 4: Character tables of D_8 and Q_8, with conjugacy classes listed as identity, central involution, and the three classes of size two. The tables and the fusion rings coincide, and the groups do not.

Representation	Character on D_8	Character on Q_8
Linear, trivial	$1,\ 1,\ 1,\ 1,\ 1$	$1,\ 1,\ 1,\ 1,\ 1$
Linear	$1,\ 1,\ 1,\ -1,\ -1$	$1,\ 1,\ 1,\ -1,\ -1$
Linear	$1,\ 1,\ -1,\ 1,\ -1$	$1,\ 1,\ -1,\ 1,\ -1$
Linear	$1,\ 1,\ -1,\ -1,\ 1$	$1,\ 1,\ -1,\ -1,\ 1$
Two-dimensional	$2,\ -2,\ 0,\ 0,\ 0$	$2,\ -2,\ 0,\ 0,\ 0$

Proposition 6.4 (Fusion rings of two groups of order eight)**.** *Let D_8 be the dihedral group of order eight and Q_8 the quaternion group. Their character tables and their representation fusion rings are isomorphic, and the groups are not isomorphic. Fusion-ring data therefore does not determine the group.*

Proof. Each group has four one-dimensional irreducible representations and one two-dimensional representation V, and the conjugacy classes have sizes $1, 1, 2, 2, 2$ in both cases. Table 4 records the character tables; they coincide after matching the central involution r^2 of D_8 with $-1 \in Q_8$ and the three classes of size two in order. In both cases the one-dimensional representations form the character group of the abelianization $\mathbb{Z}/2 \times \mathbb{Z}/2$, so their fusion is the group law of $\mathbb{Z}/2 \times \mathbb{Z}/2$ in both cases. In both cases $V \otimes V$ is the sum of the four one-dimensional representations: the character computation gives $\chi_{V\otimes V} = \chi_V^2$, which has inner product one with each linear character and zero with V. Finally $V \otimes \chi \cong V$ for every linear χ in both cases. The fusion rings are therefore isomorphic. The groups are not isomorphic: D_8 has elements of order two outside its center, while every element of Q_8 outside the center has order four. □

Proposition 6.4 is strictly weaker than Theorem 6.1 and says something different. It settles only the fusion-ring shadow: $\mathrm{Rep}(D_8)$ and $\mathrm{Rep}(Q_8)$ are not equivalent as monoidal categories, since their associator data differ, so the example does not reach the monoidal layer at all. Theorem 6.1 asserts the stronger failure for its isocategorical pairs, whose representation categories are equivalent as monoidal categories and still fail to recover the group. The symmetric layer carries what both the fusion ring and the monoidal layer lose.

7 Computational evidence

Three claims of this paper have computational evidence and none has a computational proof. The experiment files in `src/vol-04-yoneda-classification/` run the checks, and the module `Instances.hs` records the taxonomy declarations for the volume's four systems: a commutative probe system in the sense of Remark 2.6, a Boolean probe system, a finite-group representation system, and a quantum channel system in the sense of Remark 2.5. The profile data of Definition 2.2 appears in the module in the following executable shape, where the action field is the presheaf action of the definition: precomposition of a relationship along a probe morphism.

```
data RelationalProfile probe probeMor rel = RelationalProfile
  { profileProbe :: probe
  , profileRelations :: [rel]
  , profileAction :: probeMor -> rel -> rel
  }
```

The channel fragment of the module carries typed process data and computes with it: a channel is a list of Kraus operators between spaces of recorded dimensions, and the Choi matrix of Proposition 4.1 is computed from that data.

```
data KrausChannel = KrausChannel
  { kcInputDim :: Int
  , kcOutputDim :: Int
  , kcOperators :: [Matrix]
  }

data QuantumChannelSystem = QuantumChannelSystem
  { qcsChannel :: KrausChannel
  , qcsSymmetries :: String
  , qcsComposition :: String
  , qcsIsoKey :: String
  , qcsPhysicalKey :: String
  }
```

- The module `src/vol-04-yoneda-classification/Instances.hs` declares the taxonomy instances for the four systems and computes with the typed channel fragment. The family instances record the defining structural data of each family; the equivalence instances implement mathematical equivalence as equality of the recorded isomorphism data and physical equivalence as equality of the declared physical-equivalence key. The declaration of physical equivalence is stipulative in exactly the sense of the framework: the key is the relation the instance declares, and the open choice of physical equivalence is not settled by the declaration.
- The experiment `yoneda_profiles.py` evidences Corollary 3.2 and Corollary 3.4 on a finite test category: it computes profile values $\mathrm{Hom}(R, P)$, the induced actions, and the bijection $\mathrm{Hom}(A, B) \cong \mathrm{Nat}(y(A), y(B))$ on the tested hom-sets, and it computes the ultrafilter probe data of Remark 5.6.
- The experiment `choi_duality.py` evidences Proposition 4.1: it checks the Choi bijection, the positivity and trace conditions, and the composition identity on sampled completely positive maps.
- The experiment `fusion_rings.py` evidences Proposition 6.4 and the shadow case of Theorem 6.1: it computes the fusion rings of D_8 and Q_8, verifies that they are isomorphic, and verifies that the groups are not.

Table 5: Code evidence appendix: each computational claim of section 7 mapped to the file and function that evidences it in the source tree.

Computational claim	File	Function or test
Taxonomy instances for the volume's systems	`Instances.hs`	instance declarations for the four system types
Relational profile is executable data	`Instances.hs`	`yonedaProfile`
Choi matrix computable from Kraus data	`Instances.hs`	`choiMatrix` and `applyChannel`
Profile bijection on the test category	`yoneda_profiles.py`	`check_yoneda_bijection`
Distinct objects have distinct profiles	`yoneda_profiles.py`	`check_faithfulness`
Finite Stone space recovers the set	`yoneda_profiles.py`	`check_stone_finite`
Choi bijection and positivity	`choi_duality.py`	`check_choi_bijection`
Composition as index contraction	`choi_duality.py`	`check_composition`
Fusion rings of D_8 and Q_8 agree	`fusion_rings.py`	`check_fusion_ring_iso`
D_8 and Q_8 are non-isomorphic	`fusion_rings.py`	`check_groups_distinct`

Computational evidence is not proof, and none of the claims of this paper rests on these experiments. The experiments check that the formal statements behave as stated on finite instances, which is the role computational work can honestly play. Each experiment is deterministic, prints the outcome of each check, and terminates with a summary line naming the claims it evidences; the runs behind this paper check all rows of Table 5. Table 5 is the code evidence appendix: each row names a computational claim and the file and function that evidences it.

The taxonomy declarations record vocabulary: the abstract components of the probe system and the group system are opaque identifiers, so those declarations index the definitions of the taxonomy rather than encode mathematics at the type level. The typed channel fragment and the Python experiments are the computational evidence proper. No claim of machine verification is made anywhere in this volume; a later mechanization could audit the executable models against kernel-checked specifications, which would be a cross-audit of models rather than a verification of the claims.

8 Limitations

Each denial of the governing constraint is a precise limitation. They are stated below in the order the constraint lists them.

Physical equivalence. Corollary 3.4 shows that relational equivalence is isomorphism in the declared category. Physical equivalence is a chosen relation on representations, motivated by the observation that two mathematical descriptions may describe the same physics without being isomorphic, and the correct choice of that relation is open. The Yoneda embedding neither supplies the choice nor validates one. Any claim that two systems are physically equivalent because their relational profiles agree is therefore unsupported unless the morphisms in force have themselves been justified as physical equivalences.

Empirical adequacy. A mathematical structure can be consistent without describing our universe, and several mathematically distinct theories can produce identical predictions in an experimentally accessible regime. Empirical selection is a separate layer of the architecture from mathematical classification, and the Yoneda embedding lives entirely on the mathematical side. The layer is left

unformalized here by design: its mathematical content is the subject of observational equivalence, and the claim of this paper is only that relational data does not supply it.

Classification. Theorem 3.9 shows that full faithfulness of y is logically compatible with the absence of any finite-variety parametrization of isomorphism classes. The classification problem of the series asks for invariants separating objects up to a chosen equivalence; profiles separate objects up to isomorphism by a different mechanism, and the two must not be confused. Where classification content exists for a family of systems, it comes from the invariants of that family, not from representability.

Morphisms in force. All relational statements of this paper are conditional on a declared category of representations. The preservation laws and composition rules for the full morphism vocabulary of Phys are not fixed. Until they are, statements about $\mathrm{Hom}_{\mathrm{Phys}}(-, P)$ are conditional on the chosen fragment, and the one-category assumption itself is open.

Probe classes. Proposition 3.6 shows that restricted profiles determine objects only when the probe class is adequate. Every claim of determination by probes carries the probe class as a hypothesis.

The unification. The resemblance between the supermap representability theorem and the reconstruction islands is real and it is unproved as a principle. Section 9 states the problem; this paper does not assume its solution.

Evidence discipline. The computational layer of section 7 evidences statements already proved in the paper or cited from the literature; it proves nothing on its own, and no claim of the paper is promoted by computation. The taxonomy declarations of section 7 are executable models of the definitions, not verified artifacts, and no claim of machine verification is made anywhere in this volume. Should a later mechanization check these models against kernel-checked specifications, the result would be a cross-audit of executable models, not a verification of the mathematical claims.

What would change the picture is equally precise. A fixed list of preservation laws for morphisms would turn the schema into a category and make $\mathrm{Hom}_{\mathrm{Phys}}(-, P)$ a defined object rather than a family of declared fragments. A justified physical equivalence would tell one which morphisms deserve to be called identifications of systems, and would convert relational equivalence into a physical statement in the cases where the two coincide. Family-specific invariants would supply classification data where it exists, with the relational profile playing the role of a determination principle rather than a source of parameters. Each of these is a theorem to be proved or a definition to be fixed; none follows from the embedding.

9 Yoneda-principle synthesis

The open problems of this volume cluster around one question: is there a single principle behind relational representation of physical systems, covering the supermap line, the reconstruction theorems, and the formal Yoneda embedding at once? The problems below are stated in order of centrality. None is solved here.

The synthesis problem Does categorical supermap representability, which holds under channel-state duality, unify with the Gelfand, Stone, and Doplicher-Roberts reconstruction pattern into one Yoneda-principle account? A solution would exhibit one statement covering all three islands with explicit hypotheses, and would say what the probe categories have in common beyond their shape. A plausible starting point is to ask whether each island is an instance of one density or representability statement for a probe category attached to the object, with channel-state duality playing the role that characters play in the Gelfand case.

Physical content of profiles What additional structure on Phys, beyond local smallness of a declared fragment, is needed for $\mathrm{Hom}_{\mathrm{Phys}}(-, P)$ to carry physical content rather than purely mathematical content? The open status of the preservation laws is the concrete gap. Candidate answers include equipping morphisms with operational semantics, or restricting attention to fragments whose morphisms are already justified identifications.

Range of channel-state duality Which supermap settings admit a channel-state duality of the kind Theorem 4.3 requires, and how restrictive is that hypothesis in physical terms? Proposition 4.1 settles the finite-dimensional matrix case; the boundary of the hypothesis is unmapped. The applications in [1] show that the hypothesis is checkable theory by theory.

Symmetric layer in profiles Does the relational profile interact with the structure-preserving requirement in a positive way? Concretely: does the symmetric monoidal layer of a representation category appear as computable data in the presheaf $y(P)$, in a manner that recovers what Theorem 6.1 shows the monoidal layer loses? A positive answer would make relational data compatible with the constraint of section 6 rather than merely subject to it.

Status of the slogan A Yoneda lemma for physics is a slogan until it is a theorem with hypotheses. What precise statement would earn the name? A candidate must specify the category of representations, the morphisms in force, the notion of determination, and the layer of physical structure it addresses; the slogan alone supports none of these. Theorem 4.3 is the closest existing statement, and it is a theorem about processes in frameworks with duality, not about physics at large.

The first problem is the volume's headline: the two islands of concrete relational results, quantum supermaps and probe reconstruction, have never been shown to be instances of one theorem. Marking the synthesis as open is not a rhetorical hedge. A unification theorem would be a substantial result, and its absence bounds every claim this paper makes about a general relational principle.

10 Conclusion

The Yoneda embedding gives mathematical representations of physical systems a rigorous relational characterization. Objects are determined up to isomorphism by their admissible relationships with all probes; the determination is exact, and it is what the embedding establishes. In two concrete corners the relational content does more than hold formally: supermaps are representable in settings with channel-state duality, and reconstruction theorems recover objects from probe categories exactly. In exchange for these results one must state the declared category and its morphisms, declare the layers of structure an identification preserves, and accept that none of this yields physical equivalence, empirical adequacy, or a classification theorem.

The non-implication theorem of section 3 is the load-bearing negative result: representability holds in categories where no classification exists, so the classification content of any relational program must be proved separately, in tame regimes, from family-specific invariants. That is a narrow statement and an exact one. It leaves the research program of this series intact, because the program asks for restricted classification theorems and explicit obstructions rather than for a single global table, and it places relational characterization in its correct role: a determination principle whose physical and classificatory content are further theorems, each with hypotheses of its own.

The balance of the volume is therefore asymmetric in a deliberate way. On the positive side stand three exact results: the Yoneda lemma with its corollaries, the Choi correspondence and the supermap representability it supports, and the probe reconstruction theorems with their declared

hypotheses. On the negative side stand three refusals: relational equivalence is not physical equivalence, representability is not classification, and equivalence at one layer of structure is not equivalence at another. The refusals are not caveats appended to the results; they are theorems about what the results cannot deliver, and together they state the scope of the relational program with a precision that the slogan alone lacks. Within that scope, objects are determined by their relationships, and the remaining questions of the classification program are open in exactly the places section 9 names.

The title of the volume names a principle, and the sense of the name should be exact. The Yoneda classification principle is a principle about determination: objects of a declared category of representations are classified by their relational profiles up to isomorphism of that category. It is not a classification theorem in the sense of the program, and this paper has proved that it cannot be one without further input. Read in the determination sense, the principle is a theorem today; read in the classification sense, it is a research direction whose progress runs through tame families, declared equivalences, and the open synthesis of section 9.

A Fusion computation in detail

This appendix records the computation used in Proposition 6.4, so that the claim about fusion rings can be checked by hand. Write the irreducible representations of either group as $1, \varepsilon_1, \varepsilon_2, \varepsilon_3, V$, where the ε_i are the nontrivial linear characters and $\dim V = 2$. The tensor multiplicities are as follows in both cases.

Linear times linear is the group law of the character group $\mathbb{Z}/2 \times \mathbb{Z}/2$:

$$\varepsilon_i \otimes \varepsilon_j \cong \begin{cases} 1 & \text{if } i = j, \\ \varepsilon_k & \text{if } \{i, j, k\} = \{1, 2, 3\}, \end{cases}$$

and 1 is the unit. Linear times the two-dimensional representation is rigid:

$$\varepsilon_i \otimes V \cong V \qquad (i = 1, 2, 3).$$

Finally the square of the two-dimensional representation is the sum of the linear ones:

$$V \otimes V \cong 1 \oplus \varepsilon_1 \oplus \varepsilon_2 \oplus \varepsilon_3.$$

The identity follows from the character computation: χ_V takes values $2, -2, 0, 0, 0$ on the five classes, so χ_V^2 takes values $4, 4, 0, 0, 0$, whose inner product with each linear character is one and whose inner product with χ_V is zero. Dimensions agree, since $4 = 1 + 1 + 1 + 1$. The two groups give identical multiplicity data N_{ij}^k, which is the claim that their fusion rings are isomorphic. The failure of the data to distinguish the groups is then a property of the groups: element orders outside the center differ, so no ring isomorphism of fusion data can be induced by a group isomorphism.

The same computation is performed by the experiment `fusion_rings.py`, which constructs both character tables from the group presentations, computes the multiplicities by inner products, and compares the results.

References

[1] Wilson, M., Hefford, J., Hoffreumon, T. *Supermaps on generalised theories*. arXiv:2602.23865 (2026), preprint. `https://arxiv.org/abs/2602.23865`.

[2] Hefford, J., Wilson, M. *A profunctorial semantics for quantum supermaps*. Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science (2024) 1-15. DOI 10.1145/3661814.3662123. Preprint arXiv:2402.02997.

[3] Gelfand, I., Neumark, M. *On the imbedding of normed rings into the ring of operators in Hilbert space*. Matematicheskii Sbornik 12 (1943) 197-213. Reprinted in *C^*-Algebras: 1943-1993*, Contemporary Mathematics 167 (1994) 403-419. DOI 10.1090/conm/167/16.

[4] Stone, M. H. *The theory of representations for Boolean algebras*. Transactions of the American Mathematical Society 40 (1936) 37-111. DOI 10.1090/S0002-9947-1936-1501865-8.

[5] Doplicher, S., Roberts, J. E. *A new duality theory for compact groups*. Inventiones mathematicae 98 (1989) 157-218. `https://eudml.org/doc/143725`.

[6] Buchweitz, R.-O., Leuschke, G. J. *Factoring the adjoint and maximal Cohen-Macaulay modules over the generic determinant*. American Journal of Mathematics 129 (2007) 943-981. DOI 10.1353/ajm.2007.0022. Preprint arXiv:math/0505315.

[7] Mac Lane, S. *Categories for the Working Mathematician*. Graduate Texts in Mathematics, Springer, 2nd ed. (1998). DOI 10.1007/978-1-4757-4721-8.

[8] Bekkert, V. I., Drozd, Y. A. *Tame-wild dichotomy for derived categories*. arXiv:math/0310352 (2003), preprint. `https://arxiv.org/abs/math/0310352`.

[9] Choi, M.-D. *Completely positive linear maps on complex matrices*. Linear Algebra and its Applications 10 (1975) 285-290. DOI 10.1016/0024-3795(75)90075-0.

[10] Doplicher, S., Haag, R., Roberts, J. E. *Fields, observables and gauge transformations I*. Communications in Mathematical Physics 13 (1969) 1-23. DOI 10.1007/BF01645267.

[11] Etingof, P., Gelaki, S. *Isocategorical groups*. arXiv:math/0007196 (2000), preprint. `https://arxiv.org/abs/math/0007196`.