

Foundations of Structural Classification

Part 1 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 classification of mathematical representations of physical systems is a structural problem before it is a physical one, and this paper establishes the foundations that the problem requires. A representation is recorded as a five-part tuple of state space, observables, dynamics, symmetries, and composition rules. Morphisms are admitted in five named kinds, each with an explicit preservation law, and admissibility is fixed by six stated conditions that single out admissible records and admissible class specifications. Several results give the framework content. An explicit counterexample shows that the composite of an embedding and a coarse-graining satisfies no listed preservation law: the five morphism kinds are closed under self-composition and not across kinds, and the vocabulary resolves into kind-relative categories. Three restricted classes are then classified in full: finite classical systems under disjoint union are classified by cardinality with one-point systems as the unique irreducibles, finite dynamical systems decompose uniquely into connected components, and finite symmetric systems are classified by orbit type. Irreducibility is characterized relative to the composition rule; a two-qubit example identifies the additional factorization data that decomposition requires. The structural classification conjecture is stated with its five proof obligations as the theorem targets of the program. All results are stated relative to the admissibility axioms and the preservation laws fixed in the paper, in the standard axiom-relative form.

1 Introduction

Classifying the mathematical representations of physical systems is a structural question. It asks which representations exist, which of them describe the same system, and which are assembled from simpler ones. The Classification of Finite Simple Groups shows how much a completed structural classification can deliver and how much it can withhold. The classification is complete, every finite simple group belongs to a named family or is a sporadic group, and no uniform parameterization analogous to Dynkin diagrams is known [1, 2]. Practical identification runs through companion data of the ATLAS kind, order, Schur multiplier, outer automorphism group, presentations, and character tables with power maps [3]. A completed classification and a usable parameterization are separate achievements.

Three notions must be kept apart from the start. A *classification* is an exhaustive list of objects up to a stated equivalence. A *parameterization* is a uniform finite or continuous datum whose value identifies the object. A *recognition procedure* is an algorithm that decides identity from

recorded invariants. The finite simple groups deliver the first, partially the third, and not the second. Complex semisimple Lie algebras deliver all three, since Dynkin diagrams parameterize them [4]. A foundation for a classification program of physical systems has to state which of the three it pursues.

The foundations needed here are definitional. Before a classification theorem can be stated, the program needs a mathematical notion of a physical system, a notion of morphism, a notion of equivalence, and a condition that separates admissible records from arbitrary ones. The difficulty is that no universally accepted mathematical definition of every possible physical system exists. This paper does not remove that difficulty. It makes the difficulty explicit and workable: it fixes a record schema, a morphism vocabulary with preservation laws, an equivalence discipline, and a list of admissibility conditions, and it proves results about that framework rather than about any physical theory.

This paper contributes a five-component record for a representation with each component given a precise mathematical role, carried by worked classes of records that support every later construction. Preservation laws for the five named morphism kinds come with a proof that the kinds do not compose across kinds (Proposition 4.4), so the morphism vocabulary is a schema with kind-relative categories rather than a single one, and with a proof that each kind is closed under self-composition (Proposition 4.6). Four conditions single out admissible records, two further conditions single out admissible class specifications, and under a closed morphism kind the admissible records form a category (Propositions 6.4 and 6.5). Three restricted classes receive complete classification (Theorems 8.1, 8.2, and 8.4), and an explicit two-qubit computation shows that decomposition needs factorization data beyond the record (Proposition 7.6). The structural classification conjecture is stated with its five proof obligations, and a claim-status discipline keeps announced results, complete proofs, and machine-checked statements apart.

The scope of the program follows a founding distinction. Classifying mathematical representations of physical systems is a mathematical problem that can be formalized at once. Classifying physical systems themselves additionally requires a justified notion of physical equivalence and an empirical interpretation. This paper works entirely on the first side of that line. It also adopts the program's three levels of mathematical achievement, mathematical formulation, rigorous theorem, and physical validation, and it places its own results on the first two. Nothing in this paper reaches physical validation.

2 The classification problem

Completed classifications in mathematics fall into regimes, and the regimes determine what a classification program for physics can promise. Table 1 collects representative cases on which the program relies as background. The rows are established mathematics, cited here as registered background; no theorem of this paper depends on them, and the comparison with finite simple groups is the only one of them that shapes the constructions below.

Classification is sharp exactly where the identifying invariant is finite or parametric. Classification content can be entirely global, since Darboux's theorem removes all local invariants from symplectic geometry, so the classifying data of a classical system has to be topological [11]. Non-classifiability is itself a theorem rather than a difficulty, as the impossibility of classifying general type II_1 factors [7], the undecidability of Markov properties of finitely presented groups [12], the wildness of large representation categories [13], and the unrecognizability of certain four-manifolds [10] all show.

Reconstruction is better developed than classification. Gelfand duality reconstructs compact Hausdorff spaces from commutative C^* -algebras, and the Gelfand-Naimark theorem embeds every C^* -algebra into an algebra of operators [14]. Stone duality reconstructs spectra from Boolean

Table 1: Regimes of completed classification. Each row records what is classified, the identifying invariant data, and the kind of achievement the result represents.

Domain	What is classified	Identifying data	Achievement
Finite simple groups	Simple groups up to isomorphism	Family name or sporadic name	Complete, unparameterized [1, 2]
Complex semisimple Lie algebras	Algebras up to isomorphism	Dynkin diagram	Complete and parameterized [4]
Finite-dimensional C^* -algebras	Unital algebras up to *-isomorphism	Dimension vector of matrix blocks [5]	Complete and parameterized
Amenable factors	Hyperfinites factors of all types	Murray-von Neumann type and Connes invariants [6]	Complete in the amenable class
General type II_1 factors	Nothing beyond coarse strata	No countable invariant set	Impossible by theorem [7]
Principal G -bundles	Bundles over a paracompact base up to bundle isomorphism	Homotopy classes of maps into the classifying space [8]	Complete
Closed 3-manifolds	Manifolds up to homeomorphism	Decomposition into geometric pieces	Complete by geometrization [9]
4-manifolds	Restricted classes only	No general invariant set	Partly unrecognizable [10]

algebras [15]. The Doplicher-Roberts theorem reconstructs a compact supergroup from symmetric tensor categories under explicit hypotheses [16], and Deligne’s theorem identifies characteristic-zero symmetric tensor categories of subexponential growth as supergroup representation categories [17]. These are results about recovering objects from probes, and they are proof templates for the series. They are not classification theorems, and they do not make classification unnecessary.

Classification and reconstruction both rest on explicit mathematical primitives. Classification demands an object class, an equivalence relation, and a composable morphism vocabulary; reconstruction demands structured probes and preservation laws. This paper fixes those primitives and proves results about them, and it separates the representation of a system from any empirical realization of it.

This paper depends on no theorem of the series and on no classification theorem beyond the architectural comparison with finite simple groups. The classification theorems of Table 1 are registered background, and the conjecture below enters later work as a declared hypothesis. This paper asserts no universal classification; that statement is the conjecture’s subject. These rules make the definitions reusable: a definitional root that smuggled in results would make every dependent statement rest on them.

3 Mathematical framework

A representation of a physical system is a piece of mathematical data, and the schema below records that data in five components. The schema is a stipulation. It is proposed as a working definition, it is the definitional root of the series, and it is not claimed to capture every system anyone would call

physical.

Definition 3.1 (System record). A system record is a tuple

$$P = (X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{C})$$

with the following components. The component X is the state space or geometric structure of P . The component $\mathcal{O}$ is the algebra or collection of observables of P . The component $\mathcal{D}$ holds the evolution laws of P , given as Hamiltonian or vector-field data, an operator semigroup, a flow, or an evolution functor. The component $\mathcal{S}$ holds the symmetry and constraint data of P , the group or groupoid actions, gauge constraints, and constraint surfaces. The component $\mathcal{C}$ declares how P is to be combined with other systems, through tensor products, coupling terms, or interfacing rules.

Each component is a mathematical object of some ambient category, and the schema does not fix which. The state space may be a set, a symplectic or smooth manifold, a Hilbert space, or a lattice configuration. The observables may be real functions on a state space, a C^* -algebra, a von Neumann algebra, or a sheaf of observable collections. The composition component is the load-bearing one for the later sections: it fixes what counts as a decomposition, so irreducibility is always relative to $\mathcal{C}$. The five components need not be independent, and the record is a candidate definition rather than a theorem. Table 2 fixes the notation of the paper.

The composition component needs a typing statement, since composition is an operation on several systems at once. The component $\mathcal{C}$ of a record is a declaration: it names the construction under which the record is to be combined with other records of its class. The construction itself is class-level data, a binary operation on records with the induced transport on the remaining components. Two records combine exactly when they declare the same construction, and the result of the combination is a record of the same class. In the restricted classes below the declaration is the choice of disjoint union or of product, and the operation acts on state spaces with the induced maps on observables, dynamics, and symmetry data.

Table 2: Notation of the paper. Each symbol names one piece of the record schema or one relation on records.

Symbol	Meaning
P	A system record of Definition 3.1: a tuple of state space X , observables $\mathcal{O}$, dynamics $\mathcal{D}$, symmetries $\mathcal{S}$, and composition data $\mathcal{C}$
Phys_{fin}	The finite classical records of Definition 3.6
Phys_{dyn}	The finite dynamical records of Definition 3.7
Phys_{qub}	The bipartite qubit records of Definition 3.8
Phys_{adm}	The admissible subcategory of Definition 6.3
$\cong$	Mathematical equivalence of records, Definition 5.1
$\simeq$	The declared physical equivalence, Definition 5.2
$\mathfrak{I}_{\mathcal{C}}$	The class of records irreducible relative to $\mathcal{C}$
$\langle \mathfrak{F} \rangle_{\mathfrak{C}}$	Closure of a family collection under the constructions $\mathfrak{C}$
$\mathbf{1}_{\mathcal{C}}$	The neutral record of the composition rule $\mathcal{C}$

Remark 3.2 (Scope of the schema). The schema records representations, not the systems they represent, and it treats gauge orbits, superselection sectors, and state spaces with constraints at the level of the components that carry them rather than resolving them into separate object types. Two records may describe one physical system at different levels of coarse-graining, and the schema does not decide when. The remark is a boundary of Definition 3.1, not a defect to be patched silently in later sections.

Example 3.3 (Classical Hamiltonian record). Let X be a symplectic manifold, let $\mathcal{O} = C^\infty(X)$, let $\mathcal{D}$ be the Hamiltonian flow of a chosen Hamiltonian $H \in C^\infty(X)$, let $\mathcal{S}$ be the group of symplectomorphisms preserving H , and let $\mathcal{C}$ be the product of symplectic manifolds with coupled Hamiltonians. Poisson manifolds give the wider family, whose local models are supplied by Weinstein splitting [18].

Example 3.4 (Quantum record). Let X be a Hilbert space, let $\mathcal{O} = B(X)$ or a C^* -subalgebra of it, let $\mathcal{D}$ be a strongly continuous one-parameter unitary group or a semigroup of completely positive maps, let $\mathcal{S}$ be the unitary group commuting with $\mathcal{D}$, and let $\mathcal{C}$ be the Hilbert-space tensor product with tensor-product dynamics.

Example 3.5 (Lattice record). Let X be the set of vertices of a finite lattice, let $\mathcal{O}$ be the algebra of functions on X generated by local gauge-invariant operators, let $\mathcal{D}$ be the update rule of a lattice gauge model, let $\mathcal{S}$ be the local gauge group acting on the links, and let $\mathcal{C}$ be the gluing of lattices along shared boundaries. Lattice models of gauge theory are the class where restricted theorems of the series have proved attainable and where the continuum limit remains a limiting procedure in the sense of Definition 4.1.

These classes are deliberately small. Section 8 classifies the first two in full and adds a variant of the first with symmetry data.

Definition 3.6 (Finite classical record). Let Phys_{fin} be the class of records $P = (X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{C})$ with X a finite set, $\mathcal{O} = \mathbb{R}^X$ the algebra of all real functions on X , $\mathcal{D}$ the identity dynamics, $\mathcal{S}$ the trivial symmetry group, and $\mathcal{C}$ the disjoint union of state spaces. The empty record is admitted and is the neutral record for disjoint union, the $1_{\mathcal{C}}$ of Definition 7.1. The set X ranges over finite sets in a fixed countable universe, so the class is a set rather than a proper class, and isomorphic records with different underlying sets are identified in the quotient of Theorem 8.1.

Definition 3.7 (Finite dynamical record). Let Phys_{dyn} be the class of records $P = (X, \mathcal{O}, f, \mathcal{S}, \mathcal{C})$ with X a finite set, $\mathcal{O} = \mathbb{R}^X$, dynamics given by a self-map $f: X \rightarrow X$ of discrete time, $\mathcal{S}$ trivial, and $\mathcal{C}$ the disjoint union of state spaces with the map that acts componentwise. The empty record is admitted and is the neutral record for disjoint union. An invariant subsystem of P is a nonempty subset $Y \subseteq X$ with $f(Y) \subseteq Y$, carrying the restricted record.

Definition 3.8 (Bipartite qubit record). Let Phys_{qub} be the class of records with $X = \mathbb{C}^2 \otimes \mathbb{C}^2$ carrying a chosen factorization of the four-dimensional space into two qubit factors, $\mathcal{O} = B(X)$, $\mathcal{D}$ the identity, $\mathcal{S}$ the group $U(2) \times U(2)$ of local unitaries, and $\mathcal{C}$ the tensor product of factors. The factorization is part of the record, and Section 7 shows what fails when it is left out.

These classes are records, not physical theories. They exist to make every definition of the paper checkable on explicit objects.

4 Morphisms and preservation laws

A classification problem needs morphisms as much as objects. The morphism vocabulary below names five kinds of admissible relationship, and each kind is given a preservation law: the conditions a map of records must meet to count as a morphism of that kind. Without such laws the vocabulary is a list of words.

Definition 4.1 (Morphism kind). A morphism of records is declared in exactly one of five kinds. A *coarse-graining* forgets microscopic detail while preserving macroscopic observables. An *equivalence*

morphism is invertible and represents exact equivalence of representations. An *embedding* realizes P_1 as a subsystem, subtheory, or substructure of P_2 . A *duality* is a contravariant equivalence between two record presentations: invertible as a correspondence between distinct presentations rather than as an arrow of one category. A *limiting procedure* arises from a limit or colimit process, and the named processes of the program are the thermodynamic limit, the infinite-volume limit, the semiclassical limit, and the continuum limit.

Definition 4.2 (Preservation law). A *morphism of records* $g: P_1 \rightarrow P_2$ is a tuple $(g_X, g_{\mathcal{O}}, g_{\mathcal{D}}, g_{\mathcal{S}}, g_{\mathcal{C}})$ of maps in the ambient categories of the five components, with $g_X: X_1 \rightarrow X_2$, with $g_{\mathcal{O}}: \mathcal{O}_2 \rightarrow \mathcal{O}_1$ contravariant on observables, and with the remaining components transported compatibly with their declared structure; on records carrying dynamics the tuple intertwines the dynamics. Morphisms of records compose componentwise. On point-set records whose observables are functions on the state space, the tuple is determined by the state map through precomposition, and the laws of Definition 4.3 are stated in that specialization. A preservation law for a morphism kind assigns to each declared morphism of that kind explicit conditions on the tuple, and a composition rule stating which kind, if any, the composite of two declared morphisms belongs to. A morphism is admissible when its declared preservation law is stated. The morphism schema is the collection of all admissible morphisms over all kinds.

The laws used in this paper are the following. Two of the five kinds carry structure beyond a map of records: a duality is a contravariant equivalence between record presentations, and a limiting procedure is the leg of a cone on a diagram of records. Neither is an arrow of an ordinary record category in its own right, and the schema records both as morphism kinds whose preservation law carries the extra structure; the underlying map is the part that composes. The laws are stated for point-set records, and the quantum records of Example 3.4 require the algebraic form of the same laws: there a coarse-graining is a unital $*$ -homomorphism of observable algebras, as in the partial trace on a tensor factor, and an embedding is an inclusion of algebras, rather than a surjection or an injection of points.

Definition 4.3 (Concrete preservation laws). 1. *Coarse-graining*. The state map $\hat{g}: X_1 \rightarrow X_2$ is surjective, and every observable of $\mathcal{O}_2$ pulls back along $\hat{g}$ to an observable of $\mathcal{O}_1$. Surjectivity records that every macro-state of P_2 is realized by a micro-state of P_1 .

2. *Embedding*. The state map is injective, every observable of $\mathcal{O}_1$ extends to an observable of $\mathcal{O}_2$ on the embedded states, and the dynamics of P_1 is the restriction of the dynamics of P_2 to the embedded subsystem.

3. *Equivalence morphism*. The state map is a bijection and carries $\mathcal{O}$, $\mathcal{D}$, $\mathcal{S}$, and $\mathcal{C}$ across, so the map is invertible and its inverse satisfies the same conditions.

4. *Duality*. The declared morphism is the underlying map of a contravariant equivalence between two record presentations. The law names the two presentations, the direction of the correspondence, and the layers of structure preserved in the sense of Definition 5.3. Invertibility in an opposite category adds nothing to invertibility, since the isomorphisms of a category and of its opposite coincide; the content of the law is the pairing of presentations and the naming of preserved layers.

5. *Limiting procedure*. The declared morphism is the leg map of a cone or cocone of a diagram of records under one of the four named limit processes, and the law carries the diagram data. The diagram is indexed by a directed parameter with no terminal object, a family growing

without bound or a parameter approaching a limiting value. A map between two fixed finite records is never a limiting procedure of this kind.

A coarse-graining of finite classical records is worth writing out once. Let $X_1 = \{1, 2, 3, 4\}$ and $X_2 = \{p, q\}$ with $\hat{g}$ given by $1 \mapsto p, 2 \mapsto q, 3 \mapsto p, 4 \mapsto q$. Every function h on X_2 pulls back to $h \circ \hat{g}$ on X_1 , so the observable condition holds, and $\hat{g}$ is surjective. The fiber over p is $\{1, 3\}$ and the fiber over q is $\{2, 4\}$: the coarse-graining identifies states within a fiber and the pulled-back observables are exactly those that are constant on fibers. This is the classical picture of a macro-state. The same construction on quantum records replaces precomposition by an algebra map, and the partial trace on a tensor factor is its standard instance.

The point-set laws above are specializations. Records whose state space is a Hilbert space carry observables that are operators rather than scalar functions, and their morphisms are algebra homomorphisms contravariant on observables and, where states are the primitive data, maps on density matrices. The five kinds and their preservation discipline survive the change of ambient category; only the form of the state map changes.

Table 3: The five morphism kinds and the preservation conditions of Definition 4.3. The composition rule of each law is separate data, and Proposition 4.4 shows that the kinds as listed do not carry one.

Kind	State-level condition	Observable-level condition
Coarse-graining	Surjective map of state spaces	Every target observable pulls back to a source observable
Embedding	Injective map of state spaces	Every source observable extends along the embedding
Equivalence morphism	Bijection of state spaces	The whole record is carried across
Duality	Contravariant equivalence between record presentations	Preserved layers are named explicitly
Limiting procedure	Leg map of a cone on a limit diagram	The index has no terminal object

The composition rule is where the vocabulary breaks. The five kinds are not closed under composition, and the failure is visible on the smallest records in the paper.

Proposition 4.4 (Composition is not typed). *There are records P_1, P_2, P_3 and morphisms $e: P_1 \rightarrow P_2$ of embedding kind and $g: P_2 \rightarrow P_3$ of coarse-graining kind whose composite satisfies none of the five preservation laws of Definition 4.3.*

Proof. Take finite classical records with $X_1 = \{a, b\}$, $X_2 = \{1, 2, 3, 4\}$, $X_3 = \{p, q\}$. Let e have state map $a \mapsto 1$ and $b \mapsto 3$, and let g have state map $1 \mapsto p, 2 \mapsto q, 3 \mapsto p, 4 \mapsto q$. Both are admissible: e is injective and every function on X_1 extends to X_2 , and g is surjective and every function on X_3 pulls back to X_2 . The composite $g \circ e$ is the constant state map $a \mapsto p$ and $b \mapsto p$. The constant map is not surjective, so it fails the coarse-graining law. It is not injective, so it fails the embedding law. It is not bijective, so it fails the equivalence law, and it is not invertible in any pairing of record presentations, so it fails the duality law. It is not a limiting procedure, since a map between two fixed finite records is excluded from that law by Definition 4.3. The composite therefore satisfies no law in the list. $\square$

Corollary 4.5 (Schema, not category). *The morphism schema of Definition 4.2 is not a category. Turning it into one requires one of three decisions: enlarge the kind vocabulary so that composites are retyped, restrict admissible morphisms to those whose composites stay inside the list, or replace the*

single collection by a multi-category whose composition is typed in each slot. Whether one category suffices is the question these three decisions address.

The composition failure constrains every later statement about morphisms. Any classification theorem that speaks about morphisms of Phys has to state which of the three decisions it has taken, and any theorem about composition of coarse-grainings, embeddings, or dualities is conditional on that choice.

The obstruction concerns composites of different kinds. Within a single kind the laws are closed.

Proposition 4.6 (Self-composition within a kind). *Coarse-grainings compose to coarse-grainings, embeddings compose to embeddings, and equivalence morphisms compose to equivalence morphisms, under the laws of Definition 4.3.*

Proof. Surjective state maps compose to surjective maps, and pullbacks along a composite are the composite of the pullbacks, so a composite of coarse-grainings meets the coarse-graining law. Injective state maps compose to injective maps, and extensions along a composite are the composite of the extensions, so a composite of embeddings meets the embedding law. Bijections carrying the record across compose to bijections carrying the record across. Duality and limiting procedures require the extra data of their laws, and nothing here asserts closure for them. $\square$

Proposition 4.4 and Proposition 4.6 together locate the difficulty: typed composition holds inside each kind and fails across kinds. Physical modeling composes kinds freely, an embedding followed by a coarse-graining being the common case, so the single-kind restriction of Corollary 4.5 is a working convention of stated scope: admissible morphisms are restricted to one kind at a time, and Section 6 builds the resulting categories. A native treatment of inter-kind composition takes a bicategory of spans or cospans, or a typed multicategory, in which a composite carries its two kinds as typing data; that construction is the route beyond the present convention.

5 Equivalence and structural discipline

The classification problem is a quotient problem: classify objects up to some equivalence. Which equivalence is chosen decides what the classification means, and the two notions the program needs are different in strength.

Definition 5.1 (Mathematical equivalence). Two records P_1 and P_2 are mathematically equivalent, written $P_1 \cong P_2$, when there is a bijection of state spaces carrying $\mathcal{O}$, $\mathcal{D}$, $\mathcal{S}$, and $\mathcal{C}$ across in the ambient category of each component. Mathematical equivalence is isomorphism of records.

Definition 5.2 (Physical equivalence). Physical equivalence $\simeq$ is a chosen equivalence relation on records, declared per class of systems, that identifies representations describing the same physical system. Two mathematical descriptions may be physically equivalent without being mathematically isomorphic.

The second definition is a stipulation, and its content is that the choice carries weight. Dual theories can describe the same observable physics through different mathematical structures, so a classification that uses isomorphism in place of $\simeq$ classifies representations rather than systems. The program therefore requires that every use of $\simeq$ be declared, and Section 8 declares it explicitly for the restricted classes it classifies.

Definition 5.3 (Preserved-structure record). An identification of two records must name the layers of structure it preserves. The layers are the symmetric monoidal layer, the monoidal layer without

symmetry, the braided layer, the plain tensor layer, the $*$ -structure of operator-algebraic data, the algebraic structure, the topological structure, complete positivity of channels, and the unitary freedom of dilations. An identification that names no preserved layer is not admissible.

The requirement is a discipline with a proved reason. Tensor-category equivalence that discards the symmetric layer does not determine the group: non-isomorphic finite groups can carry equivalent plain tensor representation categories [19]. Any classification built on an equivalence that forgets the monoidal or symmetric layer is therefore false as a classification of the underlying structures, and every equivalence record in the series must state what it preserves.

Remark 5.4 (The Yoneda qualification). The Yoneda embedding $P \mapsto \text{Hom}(-, P)$ is established mathematics: in a locally small category an object is determined up to isomorphism by its representable functor [20]. The physical interpretation of that embedding is the subject of the following constraint, adopted verbatim. “Yoneda does not, by itself, establish physical equivalence, empirical adequacy, or a classification theorem.” Relational data alone settles none of the three, and a later volume that derives any of the three from representability alone has derived nothing.

Physical equivalence is also distinct from empirical indistinguishability, and the program keeps the two apart by construction.

Definition 5.5 (Three layers). The program separates three layers of achievement. The mathematical layer classifies admissible structures up to the chosen equivalence. The realization layer assigns to classified structures physical systems. The selection layer assigns to physical systems the observational constraints or predictions that experiments bear on. A complete solution of the first layer would not by itself solve the second or the third.

The separation has two consequences that later volumes inherit. A mathematical structure can be consistent without describing the observed world, and several mathematically distinct theories can produce identical predictions in an experimentally accessible regime. The layers belong to the assembly stage of the series, and the present paper works on the first of them only.

Remark 5.6 (Observational equivalence). A third relation is often confused with the first two. Two records are observationally equivalent when no measurement available to a stated observer distinguishes them, which is a relation on records relative to a measurement class and an observer, and not an intrinsic relation on mathematical structures. The implication pattern is one-way in each direction that matters: mathematical equivalence implies observational equivalence for any fixed measurement class, while observational equivalence does not imply mathematical equivalence, and physical equivalence is asserted on physical grounds rather than derived from either. Identifying mathematical isomorphism with empirical indistinguishability is therefore a category error, and no classification statement in this series is allowed to make that identification.

6 Admissible physical systems

Arbitrary records are not physical systems. A record with contradictory components, an observable collection that is not closed under its own operations, or dynamics that is not defined on the whole state space is a formal tuple and nothing more. The program needs an explicit list of consistency conditions, and the list below is what this paper must fix.

Definition 6.1 (Admissible record). A record $P = (X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{C})$ is admissible when it satisfies the following four conditions.

- AC-1. *Total record.* All five components are defined, and each is an object of a stated ambient category.
- AC-2. *Observable closure.* The collection $\mathcal{O}$ is closed under the operations of its ambient algebra.
- AC-3. *Well-posed dynamics.* The dynamics $\mathcal{D}$ acts on the whole state space or observable collection, as a flow, a semigroup action, or an evolution functor.
- AC-4. *Symmetry compatibility.* The symmetries and constraints $\mathcal{S}$ act on X preserving the declared structure, and $\mathcal{O}$ is invariant or equivariant under that action.

Definition 6.2 (Admissible class specification). A class specification $(\mathcal{K}, \text{Mor}_{\mathcal{K}}, \simeq_{\mathcal{K}})$ is admissible when its records satisfy Definition 6.1 and the specification itself satisfies the following two conditions.

- AC-5. *Composition closure.* The class $\mathcal{K}$ is closed under the construction of $\mathcal{C}$ wherever that construction is defined, and the composition rule declares its domain of composability as part of $\mathcal{C}$.
- AC-6. *Declared equivalence.* The specification states which equivalence relation of Definition 5.2 is in force and which preservation laws of Definition 4.3 its morphisms obey.

Definition 6.3 (Admissible subcategory). Let Phys_{adm} be the class of records satisfying Definition 6.1, inside a class specification satisfying Definition 6.2. Statements about morphisms are made in one of the categories $\text{Phys}_{\text{adm}}^K$, whose objects are the records of Phys_{adm} and whose morphisms are those of a single kind K closed under self-composition, together with the identities; Proposition 6.4 supplies these categories. The five-kind schema is not one category, by Proposition 4.4, and a morphism statement that crosses kinds has no ambient category until one of the resolutions of Corollary 4.5 is adopted.

Table 4: The six admissibility conditions: AC-1 to AC-4 constrain a record (Definition 6.1) and AC-5 to AC-6 constrain a class specification (Definition 6.2), with the subject each condition constrains and the failure it excludes.

Condition	Subject	Failure excluded
AC-1 Total record	A record	A component left implicit or undefined
AC-2 Observable closure	A record	An observable collection not closed under its own operations
AC-3 Well-posed dynamics	A record	An evolution that does not act on the whole record
AC-4 Symmetry compatibility	A record	A symmetry action that breaks the observables or the constraints
AC-5 Composition closure	A class specification	A composition rule with an undeclared domain of composability
AC-6 Declared equivalence	A class specification	An identification with no stated equivalence or preservation law

The conditions are the axioms of the framework: strong enough to yield a category and weak enough to admit the restricted classes of Definition 3.6 through 3.8. The series states them explicitly, and results over admissible records are stated relative to them.

Proposition 6.4 (Admissible morphisms form a category). *Fix a morphism kind whose preservation law contains the identities and is closed under self-composition, and take the admissible morphisms to be those of the chosen kind. Then the admissible records of Definition 6.3 together with these morphisms form a category. For the laws of Definition 4.3 the coarse-graining kind, the embedding kind, and the equivalence kind qualify, by Proposition 4.6.*

Proof. Composition of admissible morphisms is admissible by hypothesis, and the identity on a record satisfies the identity conditions of each law. Associativity and the identity laws hold because the morphisms are maps of records and composition of maps is associative. The composition rule of Definition 4.2 is exactly the hypothesis needed, which is why the open status of that rule is the open status of the category structure. $\square$

Proposition 6.5 (Conditionality over admissible systems). *Every statement quantified over Phys_{adm} is stated relative to the chosen admissibility conditions and preservation laws. Removing a condition enlarges the class and can falsify statements that hold for the smaller class.*

Proof. Consider the sentence Φ : for every record P in the class with $\mathcal{O}(P) \subseteq \mathbb{R}^X$ and all $f, h \in \mathcal{O}(P)$, the sum $f + h$ lies in $\mathcal{O}(P)$. With AC-2 in force, Φ holds on Phys_{adm} , because AC-2 requires $\mathcal{O}$ to be closed under the operations of its ambient algebra and addition is one of them. With AC-2 dropped and the remaining conditions kept, the record with $X = \{1, 2\}$, $\mathcal{O} = \{\delta_1, \delta_2\}$ of point masses, trivial dynamics, and trivial symmetries is admitted, and Φ fails on it since $\delta_1 + \delta_2$ is the constant function 1 and is not in $\mathcal{O}$. The class of true statements therefore depends on the list. $\square$

Results quantified over admissible records are stated relative to the list chosen here, and by Proposition 6.5 a different list yields a different class of true statements.

Each condition excludes a concrete pathology, and the list can be tested record by record. A tuple with a component named but not defined fails AC-1. An observable collection such as the two point masses on a two-point state space fails AC-2. A Hamiltonian vector field defined only on an open subset of the state space fails AC-3. A symmetry action that moves an observable outside the observable collection fails AC-4. A composition rule that produces records with undefined components on some inputs, with no declared domain, fails AC-5. A class of records with no declared equivalence relation fails AC-6. The conditions are strong enough to rule these out and weak enough to admit the restricted classes of Section 8, and that pairing is the design criterion for any revision of the list.

7 Irreducibility and decomposition

Which systems are simple is not a property of a record alone. It is a property of a record together with the composition rule $\mathcal{C}$, and the same object can be irreducible under one rule and decomposable under another.

Definition 7.1 (Irreducibility). Fix a composition rule $\mathcal{C}$ and write $\mathbf{1}_{\mathcal{C}}$ for its neutral record, the unit of $\mathcal{C}$ when one exists. A record P is irreducible relative to $\mathcal{C}$ when P is not isomorphic to $\mathbf{1}_{\mathcal{C}}$ and admits no decomposition $P \cong \mathcal{F}(P_1, P_2)$ with $\mathcal{F}$ the composition construction of $\mathcal{C}$ and neither P_1 nor P_2 isomorphic to $\mathbf{1}_{\mathcal{C}}$. The class of irreducibles is $\mathfrak{J}_{\mathcal{C}}$.

Definition 7.2 (Decomposition). A decomposition of P is an expression $P \cong \mathcal{F}(I_1, \dots, I_n)$ with each I_i in $\mathfrak{J}_{\mathcal{C}}$ and $\mathcal{F}$ an admissible construction. The unit record is the empty decomposition. Uniqueness of decomposition holds in the restricted subclasses below and is not claimed in general.

Definition 7.3 (Closure). For a collection $\mathfrak{F}$ of records and a set $\mathfrak{C}$ of admissible constructions, the closure $\langle \mathfrak{F} \rangle_{\mathfrak{C}}$ is the smallest class of records containing $\mathfrak{F}$ and closed under the constructions of $\mathfrak{C}$. The symbol $\mathcal{C}$ stays reserved for the composition coordinate of an individual record, and the set $\mathfrak{C}$ is program-level data.

The constructions of the program are the tensor product, the direct sum, extension, limits and colimits, and factorization of a representation into factors, and each use of a construction must state the composition data it contributes to $\mathcal{C}$. In the restricted classes below only two of these are active: disjoint union, which is a direct sum of state spaces, and the cartesian product, which is a tensor product of finite sets. The restriction is deliberate. A framework that cannot state and prove a theorem about its simplest constructions has no standing to state one about tensor products of quantum field theories.

The dependence of irreducibility on $\mathcal{C}$ is checkable on the finite classical records in two ways.

Example 7.4 (Irreducibles under disjoint union). In Phys_{fin} with $\mathcal{C}$ the disjoint union of state spaces, the neutral record is the empty record, and Definition 7.1 excludes it from the irreducibles. The irreducibles are exactly the one-point records. A record with more than one point splits as $X = \{x\} \cup (X \setminus \{x\})$ into two nonempty records, and a one-point record admits no split into two nonempty parts.

Example 7.5 (Irreducibles under product). Take the same records with $\mathcal{C}$ the cartesian product of state spaces and product composition rules. The neutral record is the one-point record, and it is the empty product. The irreducibles are the records of prime cardinality. A record of cardinality n decomposes as a product of records of prime cardinality, and the multiset of prime factors is unique up to order by unique prime factorization.

The two examples classify the same class of records and disagree on the irreducibles. Irreducibility is relative to $\mathcal{C}$, and no statement about simple systems is meaningful until the composition rule is part of the record.

Uniqueness of decomposition fails as soon as the record carries less than the composition data. The smallest instance is two qubits.

Proposition 7.6 (Factorization is extra data). *For any four-dimensional complex Hilbert space H there exist two factorizations $H \cong \mathbb{C}^2 \otimes \mathbb{C}^2$ and a unit vector $\psi \in H$ whose entanglement entropy is 0 in one factorization and $\log 2$ in the other. The bare record $(H, B(H), \text{id}_H, U(H))$ does not determine the decomposition of H into factors.*

Proof. Write $v_1 = e_1 \otimes e_1$ and $v_2 = e_2 \otimes e_2$. In the standard factorization $\psi = v_1$ is a product vector, its Schmidt rank is one, and its entanglement entropy is 0. Let W be the unitary of H that acts on the span of v_1, v_2 by the rotation $Wv_1 = (v_1 + v_2)/\sqrt{2}$ and $Wv_2 = (v_1 - v_2)/\sqrt{2}$, and acts as the identity on the orthogonal complement. Then W is self-adjoint and unitary on that span, so W^2 is the identity there. Write $\Phi_1: H \rightarrow \mathbb{C}^2 \otimes \mathbb{C}^2$ for the standard factorization, so that $\psi = v_1$ has coordinate vector v_1 . The second factorization is the pullback $\Phi_2 = \Phi_1 \circ W^{-1}$: the vector ψ is unchanged and the factorizing isomorphism changes. In the coordinates of Φ_2 the fixed vector ψ is represented by $\Phi_2(\psi) = \Phi_1(W^{-1}\psi)$, and since W^2 is the identity on the span of v_1, v_2 we have $W^{-1} = W$ there, so the coordinate vector is $W\psi = (v_1 + v_2)/\sqrt{2} = (e_1 \otimes e_1 + e_2 \otimes e_2)/\sqrt{2}$. Its Schmidt coefficients are $1/\sqrt{2}$ and $1/\sqrt{2}$ and its entanglement entropy is $\log 2$. Both factorizations are legitimate structures on H , so the decomposition into factors is not a function of the bare Hilbert space. $\square$

Two qualifications bound the claim. Factorizations related by a local unitary or by the swap of the two factors give the same entanglement entropy for every vector, so the proposition asserts the existence of a differing pair of factorizations and does not claim that every pair differs. When the symmetry data $\mathcal{S}$ of the record is fixed as a concrete subgroup $U(2) \times U(2)$ of $U(4)$, as in Definition 3.8, that subgroup does pick out a tensor factorization up to factor swap, which is why Definition 3.8 makes the factorization part of the record.

Entanglement is a relation between a vector and a chosen tensor structure, and a classification of quantum systems that forgets the tensor structure classifies vectors in a bare space. The non-uniqueness visible here is the smallest member of the family of phenomena, entanglement, emergence, and gauge redundancy among them, that make decomposition uniqueness open for interacting systems in general.

8 Classification in restricted classes

A framework for classification is worth having when classification theorems can be stated and proved inside it. Two restricted classes admit complete classification, and the theorems below give the series its smallest worked instances: an irreducible list, a decomposition theorem, a complete invariant, and a proof.

Theorem 8.1 (Finite classical classification). *Let P and Q be records of Phys_{fin} with state spaces of cardinality m and n , with composition rule the disjoint union, and with mathematical equivalence as the declared equivalence relation. Then:*

1. P is irreducible if and only if $m = 1$;
2. P decomposes as a disjoint union of m one-point records, and this decomposition is unique up to order and isomorphism of factors;
3. $P \cong Q$ if and only if $m = n$;
4. the cardinality $|X|$ is a complete invariant, and the quotient $\text{Ob}(\text{Phys}_{\text{fin}})/\cong$ is in bijection with the nonnegative integers, the empty record corresponding to zero.

Proof. Claim 1 is Example 7.4. For claim 2, the state space is the disjoint union of its points, each point carries a one-point record, and cardinality is additive under disjoint union, so any decomposition into irreducibles has exactly m one-point factors. For claim 3, an isomorphism of records in this class is a bijection of state spaces, since the observable algebra is the full algebra of functions, the dynamics is the identity, and the symmetry data is trivial; finite sets are in bijection exactly when their cardinalities agree. Claim 4 restates claim 3 as the classification of the quotient. $\square$

The classification is complete and parameterized, in the regime of Table 1 where classifications are sharp. It is also close to trivial, which is the point: the framework supports a full theorem with a proof, and the theorem's simplicity is measured evidence that the definitions are well formed rather than evidence about physics.

Theorem 8.2 (Finite dynamical decomposition). *Let P be a record of Phys_{dyn} with state space X and dynamics $f: X \rightarrow X$, and call two states adjacent when one is the image of the other under f . Then:*

1. P is irreducible if and only if the functional graph of f is weakly connected;

2. P decomposes as the disjoint union of the records on its connected components, and this decomposition is unique up to order and isomorphism of factors;
3. $P \cong Q$ if and only if the multisets of connected-component isomorphism classes of P and Q agree.

Proof. A decomposition $P \cong P_1 \sqcup P_2$ corresponds to a partition $X = X_1 \cup X_2$ with $f(X_i) \subseteq X_i$ for both parts. If the functional graph is disconnected, the parts of a disconnection are invariant and give a decomposition into two non-unit records. If the graph is weakly connected and such a partition existed, then for every edge $(x, f(x))$ the point x lies in some part X_i and invariance forces $f(x)$ into the same part, so no edge crosses the partition and the graph would be disconnected, which is a contradiction. This proves claim 1. For claim 2, every connected component is invariant, so the components give a decomposition, and each component is irreducible by claim 1. Any decomposition into irreducibles is a partition of X in which every block is invariant, and the argument just given shows that no edge crosses between blocks, so every block is a union of connected components; irreducibility of a block forces it to be exactly one component. The multiset of factors is therefore the multiset of components and does not depend on the decomposition. For claim 3, an isomorphism of records is a bijection $\varphi: X \rightarrow Y$ with $\varphi \circ f = g \circ \varphi$, and such a bijection carries components bijectively to components, preserving the isomorphism class of each. Conversely, isomorphisms between matched components patch to an isomorphism of records. $\square$

Irreducibility in Definition 7.1 is indecomposability under $\mathcal{C}$, and it is weaker than the absence of invariant subsystems. On $X = \{1, 2\}$ with $f(1) = 1$ and $f(2) = 1$, the functional graph is weakly connected, so the record is irreducible, while $Y = \{1\}$ is a proper invariant subsystem whose complement $\{2\}$ is not invariant. The distinction runs parallel to the one between simple and indecomposable modules when Maschke's theorem fails, and it keeps claim 2 of Theorem 8.2 from being read as a statement about subobjects.

The second theorem is complete without being parameterized. The classifying datum of a connected component is the isomorphism class of its functional graph, a directed cycle with rooted trees attached to its vertices, and no single integer records that datum. The class is decidable and enumerable, and the classification runs through graph isomorphism of the components rather than through a finite invariant list. The regime distinction of Table 1 reappears inside one restricted class.

The theorems of this section are statements of the combinatorial layer. The observable algebra in these classes is the full function algebra, and the interaction of nontrivial observables, Hamiltonian dynamics, and gauge constraints under morphisms is the subject of later work. On records where every claim is decidable, the definitions of irreducibility, decomposition, generation, and classification compose correctly and are proved in full.

Both theorems declare their equivalence relation explicitly, as condition AC-6 requires: the declared relation is mathematical equivalence of records. A different declared relation would give a different quotient, and the program records the choice rather than assuming it.

Theorem 8.3 (Generation in the restricted classes). *Let P be a record of one of the two restricted classes.*

1. *If P is in Phys_{fin} , then P is a disjoint union of one-point records, and under the product rule P is a product of prime-cardinality records.*
2. *If P is in Phys_{dyn} , then P is a disjoint union of records on weakly connected functional graphs.*

3. *In both cases the generation is effective: the irreducible factors are computed from the state space and the dynamics by finite enumeration.*

Proof. Claim 1 restates Theorem 8.1 claim 2 and Example 7.5. Claim 2 restates Theorem 8.2 claim 2. Claim 3 holds because both classes are finite by definition: listing the points of X produces the one-point factors, prime factorization produces the product factors, and breadth-first search in the functional graph produces the components. $\square$

Theorem 8.3 is the structural generation statement of the program proved for two restricted classes. It settles the generation obligation in these classes and nowhere else, and it is one more instance of the pattern recorded for the program at large: generation is proved in restricted domains and open in general.

Symmetry data can be added without losing classification. The third restricted class takes a finite symmetry group acting on the state space.

Theorem 8.4 (Finite symmetric classification). *Fix a finite group G , and let $\text{Phys}_{\text{fin}}^G$ be the records of Phys_{fin} with $\mathcal{S} = G$ acting on X by permutations, with $\mathcal{O} = \mathbb{R}^X$ carrying the induced action, and with $\mathcal{C}$ the disjoint union of G -sets. Then:*

1. *P is irreducible if and only if the action of G on X is transitive;*
2. *P decomposes as the disjoint union of its orbit records, and this decomposition is unique up to order and isomorphism of factors;*
3. *$P \cong Q$ if and only if for every conjugacy class of subgroups of G the number of orbits with stabilizer in that class agrees;*
4. *the resulting orbit-counting function is a complete invariant with finite range.*

Proof. A partition of X into invariant nonempty parts exists exactly when the action is not transitive, which proves claim 1 in the same way as Theorem 8.2 claim 1. The orbits are invariant and give a decomposition into transitive records, and any decomposition into irreducibles is a partition into invariant parts, each of which is a single orbit by claim 1, so the multiset of orbits is the multiset of factors and is independent of the decomposition; this proves claim 2. Two transitive G -sets are isomorphic exactly when their stabilizer subgroups are conjugate, so an isomorphism of records matches orbits by stabilizer class, and matched orbits patch to an isomorphism of records; this proves claim 3. Claim 4 follows because a finite group has finitely many conjugacy classes of subgroups and each orbit count is bounded by $|X|$. $\square$

Theorem 8.4 sits in the parameterized regime of Table 1 and shows that symmetry does not by itself push a restricted class out of it. Its invariant is already a function on conjugacy classes rather than a single number, which is the shape that classifying data takes as soon as group actions enter.

These results do not touch the candidate structural families of Section 9, interacting systems, or any physical theory. What they establish is that the framework of Sections 3 to 6 supports complete restricted classification with proofs, which is a first success criterion of the program.

9 The structural classification conjecture

The classification problem of the program is now definable.

Definition 9.1 (Classification problem). The classification problem is $\text{Classify}(\text{Ob}(\text{Phys}_{\text{adm}})/\simeq)$, the classification of admissible records up to the declared physical equivalence. The problem is stated on objects and the declared equivalence; it does not require the five-kind morphism schema to be a category, and morphism statements enter it only through the kind-relative categories of Definition 6.3.

The program’s organizing conjecture asserts that this problem has an answer built from fundamental families and constructions. The families in question are the ten candidate organizing families named in the series ontology: symplectic and Poisson structures, Hilbert spaces and operator algebras, smooth manifolds and geometric structures, principal bundles and connections, probability and measure-theoretic structures, topological and higher-categorical structures, dynamical systems and operator semigroups, tensor categories and compositional structures, discrete combinatorial structures, and derived and homotopical structures. They are candidates, not a completed taxonomy, and some may turn out to be special cases of others.

Conjecture 9.2 (Structural classification). *There is a collection of fundamental structural families $\mathfrak{F} = \{\mathfrak{F}_\alpha\}_{\alpha \in A}$ and a set $\mathfrak{C}$ of admissible constructions such that*

$$\text{Ob}(\text{Phys}_{\text{adm}})/\simeq = \text{Ob}(\langle \mathfrak{F} \rangle_{\mathfrak{C}})/\simeq,$$

where $\langle \mathfrak{F} \rangle_{\mathfrak{C}}$ is the closure of Definition 7.3 and the equality means that the two classes of records agree up to the declared physical equivalence. The notation $\text{Phys}_{\text{adm}} \simeq \langle \mathfrak{F} \rangle_{\mathfrak{C}}$ abbreviates this equality of quotients.

The conjecture enters later work as a declared hypothesis rather than as a lemma. Its five proof obligations, recorded in Table 5, are the targets that would turn it into a theorem.

Table 5: The five proof obligations of Conjecture 9.2 and the content each one requires.

Obligation	Content	Status
Completeness	Every admissible record lies in the closure of the families	Open
Soundness	Every generated record satisfies the admissibility conditions	Open
Equivalence	Different representations of one system are identified correctly	Open
Structural characterization	The families have precise mathematical definitions	Partial
Decomposition	Targeted classes are constructed from the families	Open in general

The third obligation is constrained by the preserved-structure discipline of Definition 5.3: an equivalence that discards the symmetric layer cannot identify the systems correctly, and counterexamples are proved [19]. The fifth obligation has complete instances in restricted domains and none in general.

Each obligation has a concrete target. Completeness asks for a construction of every admissible record from family data, which in the restricted classes of Section 8 is Theorem 8.3. Soundness asks for proofs that the constructions preserve the conditions of Definitions 6.1 and 6.2, which is a question about the constructions and the conditions together. Equivalence asks for a justified choice of $\simeq$ together with proofs that the identification respects the preserved layers. Structural characterization asks for membership definitions of the ten candidates named in Section 9. Decomposition asks for existence and, in restricted subclasses, uniqueness of the factorization into family data. Each of the five obligations has restricted instances in the classes of Section 8, and the general statements are the theorem targets.

The program accepts a negative outcome as a result. A proof that a specified class admits a complete classification, and a proof that certain natural classification objectives are impossible without additional restrictions, count equally as mathematical progress. The obstruction theorems surveyed in Section 2 set the expectation: outside tame regimes a total theorem of the finite-simple-group kind is implausible, and the honest synthesis is that restricted-class classification theorems and explicit obstruction theorems together are the realistic output of the program. That synthesis is itself a conjectural expectation and is not proved here.

10 Claim-status discipline

The value of a classification program depends on the honesty of its claim labels. The series uses exactly six, and every theorem-level statement carries exactly one.

Definition 10.1 (Claim-status labels). The claim-status labels are the following. *Machine-checked*: a complete formalization has been accepted by a clean kernel-checked build, with the build archived. *Complete-informal-proof*: a complete mathematical proof exists, given in full in the record or cited from the mathematical literature. *Conditional*: the statement is proved from hypotheses recorded explicitly, and some hypothesis is itself unproved. *Conjecture*: a precise statement with no complete proof. *Computational-evidence*: support from computation that is not a proof. *Open-problem*: a stated question with neither proof nor substantial evidence.

Table 6: The six claim-status labels and what each one requires. No label may be upgraded by prose, and computational evidence never upgrades a record.

Label	Requirement
Machine-checked	A complete formalization accepted by a clean kernel-checked build
Complete-informal-proof	A complete proof, given here or cited from the literature
Conditional	A proof from hypotheses that include unproved statements
Conjecture	A precise statement with no complete proof
Computational-evidence	Computation that supports a statement without proving it
Open-problem	A question with neither proof nor substantial evidence

The discipline applies to this paper’s own statements. The propositions, corollary, and theorems of Sections 4 to 8 are proved in full above, Conjecture 9.2 is stated with its proof obligations, and the experiments of the appendix are computations in the sense of Definition 10.1.

Nothing in the program is machine-checked at the time of writing: no formalization of any record has been accepted by a clean kernel-checked build, and a formalization that is announced but withheld does not qualify for the first label. The record of the announced existence result in constructive quantum field theory reads “existence claimed; manuscript and formalization withheld; verification pending” [21, 22], and an existence result and a mass-gap result are separate claims, the Yang-Mills mass gap remaining an open problem [23].

A computation never becomes a proof, however many of them agree. The executable models of the taxonomy layer for this volume are written in Haskell and are evidence of that kind: they instantiate the definitions of Sections 3 to 7, and they are checked against the statements of Section 8 by executable tests. No such model is verified by Lean, and no claim in this paper rests on a computation rather than on a proof.

The program also separates three levels of mathematical achievement: mathematical formulation, rigorous theorem, and physical validation. The definitions of this paper sit on the first level, the

classification theorems of Section 8 sit on the second, and nothing in the program has reached the third.

11 Limitations

The scope of the framework is fixed by its preservation laws and admissibility axioms. The five morphism kinds are closed under self-composition and not across kinds (Proposition 4.4), so a theorem that composes morphisms is stated relative to a chosen repair of Corollary 4.5; whether a single category accommodates all five kinds is the question that corollary poses. Results over admissible records are stated relative to the axioms of Definition 6.1 and Definition 6.2, and Proposition 6.5 records how the class of true statements moves with the list.

The equivalence relation is declared per class and is not derived. For the restricted classes of Section 8 the declared relation is mathematical equivalence, and a physical classification of the same objects would need a different relation with its own justification. For realistic systems the choice of physical equivalence is a separate task, and the preserved-structure discipline of Definition 5.3 rules out candidate choices without supplying a positive criterion.

The classification theorems treat the restricted classes of Section 8: finite classical records, finite dynamical records, and finite symmetric records. Their scope is combinatorial; the candidate structural families, interacting systems, quantum field theory, spacetime, and gravity lie outside it. They establish that the definitions support classification statements with complete proofs.

The universal classification statement is the subject of Conjecture 9.2, whose five obligations include the finiteness of a generating family set. The literature surveyed in Section 2 is registered background of the series, and the survey that registered it searched primarily arXiv and known sources, so citations absent from this paper are not evidence of absence in the literature.

A revision of the framework has a clear path. The preservation laws need completion in three directions: the inter-kind composition rule of Corollary 4.5, the algebraic form of the laws for quantum and gauge records, and the status of dualities as equivalences between record presentations. The admissibility list needs testing against records the paper does not treat, in particular gauge records with constraints and lattice records with a continuum limit. The equivalence relation takes a positive criterion beyond the negative constraint of Definition 5.3. Each of these is a stated task with a stated failure mode.

12 Conclusion

The paper established the definitional layer the classification program needs. A representation is a five-component record, morphisms are declared in five kinds with explicit preservation laws, admissibility is fixed by four record conditions and two class-level conditions, and physical equivalence is a declared choice with a preserved-structure discipline. Inside that framework each morphism kind is closed under self-composition and cross-kind composites fall outside the five laws, so the morphism vocabulary organizes into kind-relative categories. Finite classical records are classified by cardinality, finite dynamical records decompose uniquely into connected components, finite symmetric records are classified by orbit type, and each class satisfies the generation statement from its irreducibles. A four-dimensional example shows that decomposition requires factorization data the bare record does not carry.

All results are stated relative to the admissibility axioms of Definition 6.1 and Definition 6.2, the preservation laws of Definition 4.3, and the declared equivalence of Definition 5.2. The five proof obligations of the structural classification conjecture are the theorem targets of the program, and a

negative resolution of any of them would be a result of equal value.

What follows is that the definitions are precise enough to be refutable, the classification problem is now statable, the obligations are tracked, and the restricted theorems set a standard of proof that any claim of classification in this program will have to meet.

References

- [1] D. Gorenstein, *Finite Simple Groups: An Introduction to Their Classification*, Plenum Press (1982).
- [2] R. Solomon, “A brief history of the classification of the finite simple groups,” *Bulletin of the American Mathematical Society* 38 (2001) 315–352. DOI: <https://doi.org/10.1090/S0273-0979-01-00909-0>.
- [3] J. H. Conway, R. T. Curtis, S. P. Norton, R. A. Parker, and R. A. Wilson, *ATLAS of Finite Groups*, Oxford University Press (1985).
- [4] J. E. Humphreys, *Introduction to Lie Algebras and Representation Theory*, Springer (1972). DOI: <https://doi.org/10.1007/978-1-4612-6398-2>.
- [5] J. Dixmier, *C^* -Algebras*, North-Holland (1977).
- [6] A. Connes, “Classification of injective factors,” *Annals of Mathematics* 104 (1976) 73–115. DOI: <https://doi.org/10.2307/1971057>.
- [7] D. McDuff, “Uncountably many II_1 factors,” *Annals of Mathematics* 90 (1969) 372–377. DOI: <https://doi.org/10.2307/1970730>.
- [8] J. Milnor, “Construction of universal bundles I,” *Annals of Mathematics* 63 (1956) 430–436. DOI: <https://doi.org/10.2307/1969609>.
- [9] G. P. Scott, “The geometries of 3-manifolds,” *Bulletin of the London Mathematical Society* 15 (1983) 401–487. DOI: <https://doi.org/10.1112/blms/15.5.401>.
- [10] M. Tancer, “Simpler algorithmically unrecognizable 4-manifolds,” arXiv:2310.07421 (2023). URL: <https://arxiv.org/abs/2310.07421>.
- [11] V. I. Arnold, *Mathematical Methods of Classical Mechanics*, Springer (1978). DOI: <https://doi.org/10.1007/978-1-4757-2063-1>.
- [12] M. O. Rabin, “Recursive unsolvability of group theoretic problems,” *Annals of Mathematics* 67 (1958) 172–194. DOI: <https://doi.org/10.2307/1969933>.
- [13] R.-O. Buchweitz and G. J. Leuschke, “Factoring the adjoint and maximal Cohen-Macaulay modules over the generic determinant,” arXiv:math/0505315 (2005). URL: <https://arxiv.org/abs/math/0505315>.
- [14] I. M. Gelfand and M. A. Naimark, “On the imbedding of normed rings into the ring of operators in Hilbert space,” *Matematicheskii Sbornik* 12 (1943) 197–213. Reprint DOI: <https://doi.org/10.1090/conm/167/16>.

- [15] M. H. Stone, “The theory of representations for Boolean algebras,” *Transactions of the American Mathematical Society* 40 (1936) 37–111. DOI: <https://doi.org/10.1090/S0002-9947-1936-1501865-8>.
- [16] S. Doplicher and J. E. Roberts, “A new duality theory for compact groups,” *Inventiones Mathematicae* 98 (1989) 157–218. DOI: <https://doi.org/10.1007/BF01388849>.
- [17] P. Deligne, “Catégories tensorielles,” *Moscow Mathematical Journal* 2 (2002) 227–248. DOI: <https://doi.org/10.17323/1609-4514-2002-2-2-227-248>.
- [18] A. Weinstein, “The local structure of Poisson manifolds,” *Journal of Differential Geometry* 18 (1983) 523–557. DOI: <https://doi.org/10.4310/jdg/1214437787>.
- [19] P. Etingof and S. Gelaki, “Isocategorical groups,” arXiv:math/0007196 (2000). URL: <https://arxiv.org/abs/math/0007196>.
- [20] S. Mac Lane, *Categories for the Working Mathematician*, Springer (1971). DOI: <https://doi.org/10.1007/978-1-4757-4721-8>.
- [21] OpenAI, “Building scientific understanding with academic communities,” October 2026. URL: <https://openai.com/index/building-scientific-understanding-with-academic-communities/>.
- [22] Kavli Institute for Theoretical Physics, Yang-Mills working group statement, October 2026. URL: https://www.kitp.ucsb.edu/sites/default/files/YM_working_group.pdf.
- [23] Clay Mathematics Institute, Yang-Mills millennium problem page, accessed October 2026. URL: <https://www.claymath.org/millennium/yang-mills-the-maths-gap/>.

A Code evidence

Every computational claim of this paper is listed in Table 7 with the file that carries the computation and the check that runs it. The files live under `src/vol-01-foundations/`. The Haskell module instantiates the shared taxonomy classes for the records of Definitions 3.6 to 3.8; the Python experiments enumerate the finite cases. Computational evidence is not proof. The proofs of Section 8 stand on their own, and the experiments exist to catch disagreements between the statements and their models.

One simplification in the instance layer should be recorded. The Haskell record for Definition 3.8 carries a state vector and a factorization tag, and it records the observable algebra of $B(\mathbb{C}^4)$ by dimension only; the factorization tag stores the factor dimensions and not the basis change of Proposition 7.6. The instance layer therefore demonstrates the record schema and the equivalence declarations, and the two-qubit example of Proposition 7.6 is proved in the paper rather than listed as computational evidence.

Table 7: Code evidence for the computational claims of the paper. Each row names the claim, the artifact that carries the computation, and the check performed.

Claim	Artifact	Check
Irreducibles under disjoint union are one-point records	<code>exp_disjoint_union_irreducibles.py</code>	Exhaustive enumeration up to eight points
Cardinality classifies finite classical records	<code>exp_disjoint_union_irreducibles.py</code>	Brute-force isomorphism counting between labeled records
Irreducibles under product are prime-cardinality records	<code>exp_product_irreducibles.py</code>	Factorization enumeration up to sixty points
Finite dynamical records decompose uniquely into components	<code>exp_finite_dynamics_decomposition.py</code>	Enumeration of all maps on up to four points
Composite of embedding and coarse-graining is untyped	<code>exp_morphism_composition.py</code>	Exhaustive check of the stated counterexample
Taxonomy instances compile against the shared classes and run their demo	<code>Instances.hs</code> with <code>Instances-compile-evidence.txt</code>	Build transcript and demo output recorded in the evidence file