

Fundamental Structural Families of Mathematical Physics

Part 2 of *The Classification of Mathematical Representations of Physical Systems*

Matthew Long

The YonedaAI Collaboration, YonedaAI Research Collective

Chicago, Illinois, USA

m.long@yonedalabs.ai

<https://yonedalabs.ai>

October 2026

Abstract

This paper develops the family-level taxonomy that a classification of mathematical representations of physical systems requires. Ten structural families organize mathematical structures capable of representing physical systems: 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. Each family is presented with its defining invariant, its classified islands, and its recorded obstructions. Four results govern how the taxonomy may be used. A sharpness criterion shows that a family admits a classification of a prescribed shape only when its classifying invariant takes values in a set of matching cardinality or geometric structure. A globality proposition shows that no local invariant classifies inside the symplectic family. A non-closure proposition shows that a class of objects irreducible under direct sum is not itself closed under direct sum. An independence proposition shows that two families may each carry a classifying invariant and still fail to be logically independent. The ten families are candidates for an organizing taxonomy and not a completed classification, and no completeness claim is made for them.

1 Introduction

The classification of finite simple groups supplies the architectural analogy for this program: complicated structures understood through simpler constituents, gathered into families, with a short list of exceptions recorded beside them [1, 2]. The analogy is instructive and it is also incomplete. The template program is complete and a uniform description analogous to the Dynkin diagrams is missing from it: the assessment attributed to Gorenstein that “no one has yet found a clean uniform description of the finite simple groups” is recorded in the survey literature [1, 2]. Classification and parameterization are separate achievements, and the template program has achieved the first without the second.

This paper works at the level where that distinction bites first. Before any volume can classify physical systems, it needs a vocabulary of the mathematical families capable of representing them, a record of what is classified inside each family, and a discipline for choosing the data that identification actually runs on. Those are the three deliverables here. The families are recorded as candidates, their classified islands and their obstructions are recorded with the status the literature supports,

and the invariant discipline is stated as a set of definitions and results that later volumes can instantiate.

Four results carry the argument. Theorem 3.2 bounds the shape of any classification by the shape of the classifying invariant: a family whose invariant takes values in a finite set has finitely many classes, a family with uncountably many classes admits no invariant into a countable set, and a family whose classes are the image of a morphism of finite type is parameterized by a constructible set. Theorem 3.4 shows that inside the symplectic family every local invariant is constant, so classification content there is global and topological. Theorem 16.1 shows that a class of objects irreducible under direct sum is never closed under direct sum, so passing to a closure always adds objects. Theorem 17.2 shows that two families may both carry classifying invariants and still fail to be logically independent.

Section 2 fixes notation and the record structures. Section 3 states the invariant discipline and the sharpness criterion. The ten families are treated in turn across Section 6 to Section 15. Section 16 treats simple objects and their failure to close. Section 17 treats independence and reconstructability. Section 20 records what the taxonomy cannot support. Section 21 concludes. An appendix records the computational evidence and where it lives.

2 Definitions and notation

The ontology of the series fixes the ambient vocabulary. An object of the category of representations of physical systems is a record $P = (X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{C})$ carrying a state space or geometric structure, an observable algebra or collection, dynamical laws, symmetry and constraint data, and composition or interaction rules. The five components need not be independent. The record is a candidate definition and not a theorem: a universally accepted mathematical definition of every possible physical system does not exist, and establishing one is the first major research problem of the program. Part 1 of this series fixes the vocabulary reused here.

Definition 2.1 (Morphism of records). A *morphism* $f: P_1 \rightarrow P_2$ of records is a tuple of maps, one per component, that are compatible with the stated composition rules of the two records. The classes of morphism in force below are exact equivalence (invertible), coarse-graining (forgetting microscopic detail while preserving named macroscopic observables), embedding (realizing P_1 inside P_2), duality (invertible only after passage to an appropriate opposite category), and limiting procedure. *Status: stipulative.* A single ambient category for all five classes is not established, and the preservation laws under which two morphisms of a given class compose are open. Until those laws are fixed, the record is a schema rather than a category, and every statement below names the morphism class it uses.

Two equivalence relations run through the paper. Mathematical equivalence is isomorphism of the record in the ambient category of each component. Physical equivalence, written $\simeq$, is a chosen equivalence relation on objects and is strictly coarser in general. Two mathematical descriptions may be physically equivalent without being mathematically isomorphic. Every classification statement below fixes which relation is in force.

Definition 2.2 (Family record). A *family record* is a triple $(\mathfrak{F}, V, I)$ where $\mathfrak{F}$ is a full subcategory of the category of representations, V is a set called the *invariant range*, and $I: \text{Ob}(\mathfrak{F}) \rightarrow V$ is a map that is constant on $\simeq$ -classes. The map I is a *classifying invariant* for the record when it is complete: $I(P) = I(Q)$ implies $P \simeq Q$.

Definition 2.3 (Invariant-complete sector). A family record is *invariant-complete* when its invariant is classifying in the sense of Theorem 2.2. A *sector* of a family is a full subcategory of it. A sector is invariant-complete when the restricted record is.

Definition 2.4 (Structural family). A *structural family* is a family record $(\mathfrak{F}, V, I)$ whose objects are singled out by the mathematical structure of their state space, observable algebra, or both. The families are defined by the structure each one singles out and not by a partition of records, so a record whose components carry structure of several kinds belongs to each of the corresponding families. A structural family is *fundamental* for a class of physical systems when every object in that class lies in the closure of the family under the admissible constructions. The admissible constructions are tensor product, direct sum, extension, limits and colimits, factorization, and quotient or symmetry reduction.

The closure of a collection $\mathfrak{G}$ under a set of constructions $\mathcal{K}$ is the smallest class containing $\mathfrak{G}$ and closed under $\mathcal{K}$, written $\langle \mathfrak{G} \rangle_{\mathcal{K}}$. The long-term target of the program is the identity $\text{Phys}_{\text{adm}} \simeq \langle \mathfrak{F}_1, \dots, \mathfrak{F}_n \rangle_{\mathcal{K}}$ between the class of admissible physical systems and the closure of a finite family list. That identity is a research objective. No result below assumes it.

2.1 The five components and the invariant record

The object record has five components and the invariant record has five slots. They are not the same list, and the correspondence between them is a choice that each family record has to make.

Proposition 2.5 (Component invariants are not classifying in general). *Let $P = (X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{C})$ be a record and let each of the five components carry its own invariant data. Then the tuple of the five component invariants need not be a classifying invariant for P .*

Proof. The five component invariants record each component in isolation and record nothing about how the components are coupled. Two records may therefore share all five component invariants and still differ. As a concrete case, take $\mathcal{O} = M_2(\mathbb{C})$ and $\mathcal{S} = \mathbb{Z}/2$, and let $\mathcal{S}$ act on $\mathcal{O}$ either trivially or by conjugation with $\text{diag}(1, -1)$. The algebra is the same and the group is the same, so the component invariants of $\mathcal{O}$ and $\mathcal{S}$ agree, and the two records are not isomorphic as records because the actions differ. The tuple of component invariants misses the coupling between components and cannot be classifying in general. $\square$

Theorem 2.5 is a caution about assembling a classification from parts. It is not a claim that no decomposition of a classification exists; it is a claim that the obvious decomposition of the record into five independent invariants does not produce one.

Remark 2.6 (Status discipline). Results drawn from the literature are attributed to their sources and stated with the scope their sources establish. Results proved here are proved in full. Conjectural relationships are named as conjectural. Nothing in this paper is machine checked, and no claim of machine checking is made anywhere in the series.

3 The invariant discipline

Identification in a classified domain runs on a small amount of named data. The atlas of finite groups identifies a group by its order, its Schur multiplier, its outer automorphism group, a presentation, and a character table with power maps [3]. The Killing-Cartan classification identifies a complex semisimple Lie algebra by its Dynkin diagram [4]. The highest weight theorem identifies a finite-dimensional irreducible representation by a dominant integral weight [5]. Finite-dimensional unital

complex C^* -algebras are identified by a dimension vector [9]. Principal bundles are identified by a homotopy class of maps into a classifying space together with characteristic classes [17, 18]. In every case the classifying data is finite or parametric, and classification is sharp exactly where that holds [1, 4, 5, 6, 9, 17].

Definition 3.1 (Invariant record). The *structural invariant record* of a theory T is the tuple $\mathcal{I}(T) = (\mathcal{S}_T, \mathcal{O}_T, \mathcal{D}_T, \mathcal{G}_T, \mathcal{R}_T)$ describing state invariants, observable invariants, dynamics invariants, symmetry invariants, and representation invariants. Theorem 2.5 applies to this tuple as much as it applies to the raw component data: five slots, each recording one component, still record nothing about how the components are coupled. The invariant record is therefore a catalogue of the data a classification would have to determine and not itself a classifying invariant. A classifying invariant is one of the named data types whose value determines the object up to the chosen equivalence.

The discipline has three parts. First, a family record names one classifying invariant and one invariant range, and states the equivalence the invariant is complete for. Second, the record states the scope in which completeness holds, because completeness outside scope is false. Third, when a family has no complete invariant on record, the record says so instead of inventing one. The family tables below apply this discipline row by row, and each row states its scope limit in its last column.

3.1 Sharpness of classifying invariants

Classification is sharp where the invariant is finite or parametric, and the quantitative form of that observation is the following.

Theorem 3.2 (Sharpness criterion). *Let $(\mathfrak{F}, V, I)$ be an invariant-complete family record and write $\mathfrak{F}/\simeq$ for the set of $\simeq$ -classes.*

1. *The assignment $[P] \mapsto I(P)$ is a well-defined injection $\mathfrak{F}/\simeq \hookrightarrow V$.*
2. *If V is finite with $|V| = N$, then $|\mathfrak{F}/\simeq| \leq N$.*
3. *Suppose in addition that V carries the structure of a variety over an algebraically closed field k , that W is a k -variety, and that $f: W \rightarrow V$ is a morphism of finite type whose image $f(W(k))$ contains every attained value of I . Then the attained values of I form a constructible subset of V , and $\mathfrak{F}/\simeq$ injects into it. The content of the claim is the constructibility, which is a parameterization statement; the injectivity is part (1) and is not what this part adds.*
4. *If $\mathfrak{F}/\simeq$ is uncountable, then no invariant-complete record for $\mathfrak{F}$ has countable range.*

Proof. For the first claim, I is constant on $\simeq$ -classes, so the assignment is well defined on classes, and completeness gives injectivity. The second claim is the cardinality of a subset of a finite set. For the third claim, Chevalley's theorem states that the image of a morphism of finite type between varieties over an algebraically closed field is constructible. The hypothesis places the attained values of I inside $f(W(k))$, so they form a subset of a constructible set and the classes inject into that subset by part (1). What the hypothesis buys is not the injectivity but the geometry: the classes are then described by the points of an explicitly constructible set, which is the parameterization that part (3) claims. No identification of V with the set of its k -points is being made: the invariant I is a map of sets in general, and part (3) is the case in which its range happens to carry a variety structure and its attained values happen to lie in the image of a finite-type morphism. The fourth claim is the contrapositive of the first: an injection from an uncountable set into a countable set does not exist. $\square$

Corollary 3.3 (Parameterization is a stronger achievement). *A family with a finite invariant range is classified into finitely many classes, and a family is parameterized only when its attained values form a constructible or otherwise explicitly described subset of a structured range. Classification and parameterization are separate achievements.*

Theorem 3.3 recovers the template program’s own situation in the form that applies to it. The finite simple groups have infinitely many isomorphism classes, so no finite invariant range classifies them; they are classified, and a uniform parameterization analogous to Dynkin diagrams is missing [1, 2]. Practical identification there runs through the atlas data instead [3], which is a recognition procedure in the sense of Theorem 5.4 and not a parameterization.

3.2 Locality of invariants

Classification content can be entirely global.

Proposition 3.4 (Globality of classification in the symplectic family). *Let (M, ω) and (N, η) be symplectic manifolds of the same dimension $2n$ with $n \geq 1$. Let f assign a value to every symplectic germ and be invariant under symplectomorphism of germs. Then f takes the same value at every point of M and at every point of N . Consequently no invariant defined on germs and constant on each connected component is classifying for the symplectic family at fixed dimension $2n$, and every classifying invariant for that family is global.*

Proof. By Darboux’s theorem, every symplectic germ of dimension $2n$ is symplectomorphic to the standard germ $(\mathbb{R}^{2n}, \omega_0)$ [6]. The hypothesis $n \geq 1$ is needed: at $n = 0$ the manifold is a disjoint union of points and a constant invariant does classify it, and the dimension itself is a local germ invariant once dimensions are allowed to vary. Invariance under symplectomorphism forces f to be constant on the orbit of germs, and that orbit is the whole set of $2n$ -dimensional germs. So f is constant at every point of M and of N , and it cannot distinguish them. A classifying invariant must distinguish objects that are not equivalent, so it cannot be local. $\square$

Corollary 3.5. *Classifying data for the symplectic family is topological or global in character. This reverses the Riemannian situation, in which local curvature invariants do classifying work.*

Theorem 3.4 is the reason the symplectic row of the family tables below reads “complete locally, with no local invariants”. Completeness of the local classification is a theorem, and its content is that there is nothing local to compute.

3.3 Achievement kinds for classified sectors

The recorded results fall into a small number of shapes, and distinguishing the shapes keeps the family tables from flattening a proof into a status word.

Definition 3.6 (Achievement kind). Let $\mathcal{S}$ be a sector together with a statement of what is classified and up to which equivalence. Three conditions attach to $\mathcal{S}$ and together determine its *achievement kind*. The first is whether a theorem on record names a classifying invariant for $\mathcal{S}$; this is an epistemic condition about the literature and not a claim that no set-theoretic invariant exists, since the assignment sending an object to its own equivalence class is always a classifying invariant of sets. The second is what a recorded classification classifies: existing objects, or obstructions to the existence of objects. The third, available only when the first two conditions name objects, is whether the attained values of the invariant form an explicitly described subset of a structured range. The resulting kinds are: *complete and parameterized*, *complete and unparameterized*, *complete as an*

obstruction classification, *impossible*, and *no record*. A sector is *impossible* when a theorem on record proves that no classifying invariant of the prescribed type into a set of the required cardinality exists. A sector has *no record* when no theorem names a classifying invariant and no impossibility theorem is on record.

Proposition 3.7 (The kinds are exhaustive). *For any sector $\mathcal{S}$ exactly one of the five achievement kinds of Theorem 3.6 applies, and the five are pairwise exclusive and exhaustive.*

Proof. The three conditions of Theorem 3.6 are tested in order. If no theorem on record names a classifying invariant and no impossibility theorem is on record, the kind is no record, and no other kind applies since the others all presuppose one of the two kinds of record. If an impossibility theorem is on record, the kind is impossible, and no other kind applies. If a theorem names a classification and the classified target is the obstructions, the kind is complete as an obstruction classification. If it names a classification of objects, the attained values either form an explicitly described subset of a structured range or they do not, which is exactly the remaining pair. The tests are disjoint at each stage and cover every case at each stage, so one and only one kind applies. $\square$

The five kinds are realized in the families below. The Killing-Cartan and highest weight classifications are complete and parameterized [4, 5]. The finite simple groups are classified and not uniformly parameterized [1, 2]. Anomalies are classified as obstructions to quantization [19]. General separable type II_1 factors are impossible to classify [8]. Sectors with no record at all, such as the semigroup family, realize the fifth kind, and Theorem 3.7 accounts for them rather than excluding them.

Table 1: Achievement kinds realized inside the ten families. A family occupies more than one row when its sectors differ.

Kind	Realized in	Invariant range
Complete and parameterized	Complex semisimple Lie algebras [4], finite-dimensional irreducible representations [5], finite-dimensional C^* -algebras [9]	Dynkin diagram, dominant integral weight, dimension vector
Complete and unparameterized	Finite simple groups [1, 2]	Atlas data, with no uniform description
Complete as obstruction classification	Anomalies in gauge theories [19]	Cohomological obstruction classes whose vanishing is required
Impossible	General separable type II_1 factors [8]	No range of the required cardinality exists
No record	Completely positive maps up to isomorphism [10, 11]	The normal form is not an isomorphism classification

The table makes two points that a status token would hide. A single family can occupy several kinds at once, because its sectors differ in cardinality and in what is classified. And the absence of a record is not the same as impossibility, since an impossibility is a theorem and an absence is a gap.

4 Worked family records

Three family records apply the discipline of Section 3. Each names the invariant data it supplies, the equivalence in force, and the scope in which the data is complete; where a field is absent the

record says so.

4.1 A cotangent bundle in the symplectic family

Let Q be a smooth manifold and $M = T^*Q$ with its canonical symplectic form $\omega = d\theta$, where θ is the tautological one-form. The record is: state space M , observables $C^\infty(M)$, dynamics the Hamiltonian flow of a smooth function H , symmetries the cotangent lifts of diffeomorphisms of Q preserving H , and composition the product of systems with the product symplectic form.

Local data is empty by Theorem 3.4. The equivalence in force is symplectomorphism of the total space. The record supplies a global invariant and not a classifying one, and it says so.

Proposition 4.1 (Global invariants of a cotangent bundle). *The zero section of T^*Q is a deformation retract of T^*Q . Consequently the homotopy type of the symplectic manifold T^*Q equals the homotopy type of Q , and every homotopy invariant of Q is a global invariant of the symplectic record.*

Proof. The maps $r_t: T^*Q \rightarrow T^*Q$ given by $r_t(x, \xi) = (x, t\xi)$ for $t \in [0, 1]$ form a deformation retraction: each r_t is smooth, r_1 is the identity, r_0 factors through the zero section, and $r_s \circ r_t = r_{st}$ for $s, t \in [0, 1]$. The zero section is diffeomorphic to Q , so T^*Q and Q are homotopy equivalent and share all homotopy invariants. $\square$

Theorem 4.1 supplies an invariant and not a complete one. Two homotopy-equivalent bases need not have symplectomorphic cotangent bundles, and the record also carries the Hamiltonian H , which the homotopy type does not see. The scope of the record is therefore the homotopy data of the base, and the record does not claim more.

4.2 A finite-dimensional quantum system

Take $A = M_2(\mathbb{C}) \oplus \mathbb{C}$, the observable algebra of a two-level system together with a classical bit. The record is: state space the space of block-diagonal density matrices $\rho_2 \oplus p$, observables A , dynamics a one-parameter group of $*$ -automorphisms, symmetries the unitary group of A acting by conjugation, and composition the tensor product of algebras. The state space is the state space of A and not of $M_3(\mathbb{C})$, because the states of a direct sum are pairs of states of the summands.

By Theorem 7.1 the classifying invariant of the observable algebra is the dimension vector, here the multiset $\{2, 1\}$ of matrix sizes. That invariant classifies the algebra and not the full record, which also carries dynamics and composition data, as Theorem 2.5 records. The invariant range at total complex dimension five is the set of multisets of positive integers whose squares sum to five, and there are two of them: $\{2, 1\}$ from $4 + 1$, and $\{1, 1, 1, 1, 1\}$ from five copies of 1. At total dimension four there are two classes as well, given by $\{2\}$ and $\{1, 1, 1, 1\}$. The count of classes grows with the dimension, and Theorem 3.2 bounds it by the size of the range.

4.3 Line bundles over the two-sphere

Let $G = U(1)$ and let the base be $X = S^2$. The record is: state space the space of connections on a principal $U(1)$ -bundle, observables the gauge-invariant functionals, dynamics the parallel transport of a chosen connection, symmetries the gauge group, and composition the tensor product of bundles.

By Theorem 9.1 the classifying invariant of the underlying topological bundle is the first Chern class, which takes values in $H^2(S^2; \mathbb{Z}) \cong \mathbb{Z}$. The invariant range is countably infinite and every integer is attained, so the topological sector is classified and not finite. It is an instance of the distinction in Theorem 3.3: the classification is complete, and the range is an explicit discrete group in which the attained values are described. The invariant classifies the bundle and not the space of connections on it, which is an affine space and carries no such discrete invariant.

4.4 Constructions acting on invariants

A family record does not automatically extend to the closure of its family, because the constructions move the invariant. The finite-dimensional operator-algebra sector makes this concrete.

Proposition 4.2 (Induced operations on the dimension vector). *Let A and B be unital finite-dimensional complex C^* -algebras with dimension vectors $\{n_1, \dots, n_k\}$ and $\{m_1, \dots, m_l\}$. Then the dimension vector of $A \oplus B$ is the multiset union $\{n_1, \dots, n_k, m_1, \dots, m_l\}$, and the dimension vector of $A \otimes B$ is the multiset $\{n_i m_j : 1 \leq i \leq k, 1 \leq j \leq l\}$. Consequently the range of the classifying invariant of Theorem 7.1 is closed under these two operations, and the invariant record of a composite system is computed from the records of the summands.*

Proof. The Wedderburn decomposition of a direct sum of semisimple algebras is the multiset union of the Wedderburn decompositions of the summands, so the matrix sizes add in the multiset sense. For the tensor product, $M_a(\mathbb{C}) \otimes M_b(\mathbb{C}) \cong M_{ab}(\mathbb{C})$ by the standard identification of the tensor product of matrix algebras, and tensor product distributes over direct sums in both variables. Collecting the simple summands of $A \otimes B$ gives exactly the multiset of products $n_i m_j$. Both constructions therefore compute a new classifying invariant from the old ones without reference to any other data. $\square$

Theorem 4.2 is the sense in which a classifying invariant can be functorial for the admissible constructions. It also marks a limit: the same computation in a sector where no classifying invariant is on record cannot be carried out, and the closure of such a sector is unrecorded.

5 Morphisms and invariants

The invariant discipline has to say what happens when two records are related by a morphism rather than an equivalence. The ontology distinguishes coarse-graining, equivalence, embedding, duality, and limiting procedures, and the invariant does not behave the same way along all five.

Proposition 5.1 (Invariance under equivalence only). *Let $(\mathfrak{F}, V, I)$ be a family record and let $f: P \rightarrow Q$ be an invertible morphism representing exact equivalence. Then $I(P) = I(Q)$. The conclusion fails for coarse-grainings and for embeddings in general.*

Proof. The first claim is the defining condition of Theorem 2.2: I is constant on $\simeq$ -classes, and an invertible equivalence morphism places P and Q in one class. For the negative claim, a coarse-graining forgets microscopic detail, and two systems that coarse-grain to the same macroscopic record need not be equivalent. The dimension vector of Theorem 7.1 is a classifying invariant only for the finite-dimensional algebra itself; the same observable algebra can carry inequivalent states and dynamics, so coarse-grained records identify objects that I separates. An embedding places P inside a larger system, and the invariant of the larger system is computed by Theorem 4.2 and not inherited from P . $\square$

Remark 5.2 (Dualities). A duality is generally invertible only after passage to an appropriate opposite or dual category, and it must state which structure it preserves. The Gelfand correspondence of Table 3 is an anti-equivalence, and it identifies invariant records on the two sides up to the stated opposition. A duality that does not name the preserved structure cannot be used to transfer a classification, for the reason given in Theorem 13.1.

Remark 5.3 (Limiting procedures). A limiting procedure produces a new object, and the invariant of the limit is not in general the limit of the invariants. Continuum limits are the load-bearing case:

results proved on a lattice in a restricted regime do not by themselves close the continuum limit, and the invariant record of the limiting object must be established on its own terms.

5.1 Recognition data

The atlas discipline separates three tasks that are easy to conflate [3]. Parameterization is an explicit uniform description of the classes. Classification is an exhaustive list. Recognition is an algorithm that identifies which class a given object lies in, by computing its invariant and comparing.

Proposition 5.4 (Recognition cost). *Let $(\mathfrak{F}, V, I)$ be an invariant-complete family record with V a set carrying decidable equality, and assume $I(P)$ is computable from a finite presentation of P . Recognition of the class of $P \in \mathfrak{F}$ reduces to computing $I(P)$ and comparing it with the recorded values, and it costs at most one invariant computation plus one comparison per candidate class. If in addition the list of attained values is finite and stored, recognition is an algorithm.*

Proof. Injectivity of $[P] \mapsto I(P)$ from Theorem 3.2 gives that $I(P)$ determines the class. The procedure computes $I(P)$ from the presentation of P , which the hypothesis makes possible, and looks the value up in the stored list. Decidable equality decides the comparison. The cost claim is the cost of the two steps. The computability hypothesis is doing real work: decidable equality on a finite range does not make the map I computable, and without the hypothesis the procedure is not defined on presentations at all. $\square$

Recognition is therefore the weakest of the three tasks when the invariant is computable, and it is the one the atlas supplies in practice [3]. Parameterization is the strongest. The three are separate achievements, and a family row that reports one of them does not report the others.

6 Family I: symplectic and Poisson structures

This family singles out classical mechanics. Its defining structural data is a symplectic or Poisson state space; systems in it carry function algebras of observables and Hamiltonian dynamics [7]. A record whose state space is symplectic belongs to this family whatever else its components carry.

The local picture is closed. Darboux’s theorem gives a normal form and Theorem 3.4 extracts its classification consequence. For Poisson manifolds the local picture is also closed and it is not empty: near a point of rank $2k$ a Poisson manifold splits as a symplectic factor times a transverse Poisson factor [7]. The symplectic leaves foliate the manifold, and on the dual of a Lie algebra the leaves of the Kirillov-Kostant-Souriau structure are the coadjoint orbits [7].

Theorem 6.1 (Poisson local models). *Near a point of rank $2k$ a Poisson manifold is locally isomorphic to a product of a symplectic manifold of dimension $2k$ and a Poisson manifold whose linearization vanishes at the point [39]. The symplectic leaves form a foliation, and two Poisson manifolds with isomorphic local models need not be globally isomorphic.*

The global picture is where the family’s classification content lives, and the corpus does not fix a list of global invariants. Candidates in ordinary use are the de Rham class of the symplectic form, Chern classes of compatible almost complex structures, and flux data. The periodic chart records “topology of ω ” for this row. Making that phrase precise is an open deliverable.

The link from this family to the operator-algebra family through deformation quantization is a research program and not an equivalence theorem. No result below depends on it.

7 Family II: Hilbert spaces and operator algebras

This family singles out quantum mechanics. Its defining structural data is a C^* -algebra or von Neumann algebra of observables, on a Hilbert-space state space, with tensor product as the composition rule. A record whose observable algebra is noncommutative belongs to this family whatever else its components carry.

The finite-dimensional sector is the model case for the whole paper.

Theorem 7.1 (Structure of finite-dimensional unital complex C^* -algebras). *Let A be a unital finite-dimensional complex C^* -algebra. Then $A \cong \bigoplus_{i=1}^k M_{n_i}(\mathbb{C})$ for a unique multiset $\{n_1, \dots, n_k\}$ of positive integers, and that multiset determines A up to $*$ -isomorphism. The invariant range of the sector of all such algebras is the countable set of finite multisets of positive integers. At fixed total complex dimension $d = \sum_i n_i^2$ the attained values are finite in number, and Theorem 3.2 bounds the number of classes at that dimension accordingly.*

Proof sketch. The Jacobson radical of a C^* -algebra is a closed two-sided $*$ -ideal of topologically nilpotent elements and is zero, so A is semisimple. The Artin-Wedderburn theorem then gives $A \cong \bigoplus_i M_{n_i}(D_i)$ with each D_i a division algebra over $\mathbb{C}$. Every finite-dimensional division algebra over $\mathbb{C}$ equals $\mathbb{C}$, because $\mathbb{C}$ is algebraically closed and the minimal polynomial of any element splits. The C^* -identity then determines the $*$ -structure on each matrix summand, and an isomorphism of the algebras that respects the $*$ -structure permutes the simple summands. Uniqueness of the multiset is the uniqueness of the Wedderburn decomposition of a semisimple algebra. At fixed total dimension d the number of multisets whose squares sum to d is finite, which is the stated bound. $\square$

Theorem 7.1 is the sector in which the invariant discipline behaves exactly as intended: one named invariant, a finite range, a complete determination. It is also the sector that the later volume on irreducibility and decomposition takes as its model theorem.

The infinite-dimensional picture is different in kind. Amenable factors are classified by Murray-von Neumann type and Connes invariants [8]. Beyond the amenable class, there are uncountably many pairwise non-isomorphic separable type II_1 factors, constructed by McDuff [37]. By Theorem 3.2 no invariant with countable range classifies that sector. That consequence is what uncountability alone delivers, and a stronger obstruction is known: the isomorphism relation on separable type II_1 factors is not classifiable by countable structures, in the sense of Borel reducibility [38]. The parametric question is therefore closed in the negative rather than left open. The row of the family table records the countable-invariant obstruction and cites the stronger one.

Remark 7.2 (What the impossibility does and does not say). The impossibility is a theorem of McDuff [37] about the sector of all separable type II_1 factors. It does not forbid restricted classification under extra hypotheses such as amenability, and it does not bear on the finite-dimensional sector of Theorem 7.1. The row of the family table for this family records all three situations at once, which is why a single status token per family is not enough.

Completely positive maps and quantum channels sit in this family through Stinespring dilation and the Kraus form [10, 11]. The classification status of that data is treated in Section 12, where the maps are the objects rather than the ambient structure.

8 Family III: smooth manifolds and geometric structures

This family singles out relativity. Its defining structural data is a smooth manifold with additional geometric structure, carrying function or operator algebras of observables and a flow or field equation as dynamics.

Three classification results are on record and each holds in a restricted regime. The holonomy groups of connected Lorentzian manifolds of signature $(1, n - 1)$ are classified by a Berger-list analogue, with realizing metrics constructed [27]. The causal hierarchy of spacetimes is complete as a ladder of conditions from the weakest to global hyperbolicity [26]. Closed three-manifolds are classified up to homeomorphism by geometrization: each decomposes into pieces carrying exactly one of eight Thurston geometries [35].

Dimension four is the boundary. Only restricted classes of four-manifolds are classified, by eighteen model geometries together with one countably infinite family, and recognition is not algorithmic: the homeomorphism problem for a connected sum of nine copies of $S^2 \times S^2$ is undecidable [30].

Proposition 8.1 (Recognition without classification). *There exist four-manifolds for which homeomorphism is undecidable. Consequently no algorithm whose output is a classifying invariant in a decidable range can classify all four-manifolds, and any classification statement in this regime must restrict the class of manifolds under consideration.*

Proof. Undecidability of the homeomorphism problem for the stated family of four-manifolds is proved in [30], and the analogous undecidability for group-theoretic Markov properties is the classical Adyan-Rabin theorem [29]. An algorithm producing a decidable classifying invariant would decide homeomorphism by comparing outputs, which contradicts undecidability. So no such algorithm exists on the unrestricted class. $\square$

9 Family IV: principal bundles and connections

This family represents gauge theories. Its state spaces are spaces of connections on a principal bundle, its observables are gauge-invariant functionals or operator algebras, and its symmetries are gauge group actions.

This family is the one in which the invariant discipline is closest to its classical source.

Theorem 9.1 (Classification of principal bundles). *Let X be a paracompact Hausdorff space and G a topological group. The isomorphism classes of principal G -bundles over X are in bijection with the homotopy classes of maps $X \rightarrow BG$, where BG is a classifying space for G . Characteristic classes of these bundles are exactly the pullbacks of universal classes in $H^*(BG)$, including Stiefel-Whitney, Pontryagin, Euler, and Chern classes. The Dixmier-Douady class is the corresponding invariant for $U(1)$ -bundle gerbes, which lie in $H^3(X; \mathbb{Z})$, and it classifies those rather than ordinary principal bundles.*

The invariant range here is the homotopy set $[X, BG]$ with its characteristic-class data, and Theorem 9.1 states that the invariant is complete for this sector, so the record is invariant-complete in the sense of Theorem 2.3 for every base and every group to which the theorem applies. The shape of the range varies: it is finite for some bases and groups and infinite for others, and Theorem 3.2 bounds the number of classes by the size of the range in each case.

Anomalies in this family are classified as obstructions rather than as objects. The fermionic anomaly is the non-trivializability of the determinant or Pfaffian line of the Dirac operator, and the gauge anomaly is a cohomological obstruction; quantization exists only when the relevant obstructions vanish [19]. A classification of obstructions is a complete classification of the cases in which the theory exists.

10 Family V: probability and measure-theoretic structures

This family represents statistical physics. Its state spaces are spaces of probability measures or density operators, its observables are measurable functions or operators, and its composition rule is product measure or tensor product.

The classical row carries the one uniqueness result in the ten families.

Theorem 10.1 (Chentsov). *On the classical statistical state space of a finite sample space, the Fisher metric is the unique Riemannian metric, up to a positive constant factor, that is monotone under stochastic maps. The uniqueness is recorded in the survey literature as the classical anchor of the monotone-metric classification [28], and the quantum analogue there is parametric rather than unique.*

The quantum analogue fails to be unique. Monotone metrics on quantum state spaces are classified parametrically by operator monotone functions, which form a continuum and not a finite list [28]. Classical uniqueness therefore does not transfer to the quantum case, and any classification statement in this family must record which side of that line it stands on.

Remark 10.2 (Scope of the uniqueness). The uniqueness is a statement about metrics monotone under stochastic maps on the classical state space. It is not a classification of classical statistical systems, and it is not a statement about quantum state spaces. Uniqueness and classification are separate achievements, as Theorem 3.3 records.

10.1 A binary statistical system

Take the state space of probability vectors $(p, 1 - p)$ with $p \in (0, 1)$ and the observable algebra of real functions on $\{0, 1\}$. The Fisher metric in the coordinate p is

$$g_{pp} = \frac{1}{p(1 - p)}.$$

Up to a positive constant factor this is the only Riemannian metric on this state space that is monotone under stochastic maps, by Theorem 10.1. What is unique is the metric as ambient geometry on the state space, and not the physical system record. The record carries a state space, a two-point function algebra of observables, and a Markov kernel as dynamics, and the metric does not classify the record: two binary systems with different kernels are not equivalent, and both carry the same metric up to scale. This sector therefore has no classifying invariant on record, and the uniqueness of the metric is a statement about geometry rather than about classification.

The quantum analogue of the same two-point system has a state space of 2×2 density matrices and a continuum of monotone metrics, parametrized by operator monotone functions [28]. The classical row is therefore a uniqueness statement of a kind the quantum row does not have.

11 Family VI: topological and higher-categorical structures

This family represents topological quantum field theories. Its state spaces are spaces of fields on a manifold, its observables are topological invariants, and its composition rule is gluing of manifolds.

Functorial topological quantum field theory is tractable under the Atiyah axioms, in contrast to the general non-topological case [22]. The extended case is where a classification claim exists and where the proof status matters.

Remark 11.1 (The cobordism claim). It is claimed that n -dimensional unitary extended topological quantum field theories are weak n -functors out of the free stable weak n -category with duals, and that they are classified by fully dualizable objects [20, 21]. The 2009 proof is described by its author as an informal sketch of a proof [21]. The claim is recorded here as a claim and not as a theorem, and it is not placed in a theorem environment for that reason.

Remark 11.2 (Proof gap). The classification claim of Theorem 11.1 is recorded as claimed and not as complete. Later low-dimensional cases have been proved in the literature, and which cases are complete is a matter for verification case by case. A family row that reads “claimed” is a family row that should not carry load-bearing dependencies.

12 Family VII: dynamical systems and operator semigroups

This family represents dissipative and open systems. Its state spaces are spaces of states, its observables are operators or algebras, and its dynamics is a semigroup of maps or a generator.

No classification theorem for this family is on record in the surveyed literature. What is on record is a normal form.

Theorem 12.1 (Stinespring). *Every completely positive map admits a minimal Stinespring dilation, unique up to unitary equivalence of the dilating space [40]. The associated Kraus form is a normal form, and the number of Kraus operators in a minimal dilation is an invariant of the dilation.*

The next proposition is the precise sense in which that invariant falls short.

Proposition 12.2 (Kraus rank is not injective on channels). *Fix the equivalence relation on channels to be equality of completely positive maps. Let $\Phi_U: M_2(\mathbb{C}) \rightarrow M_2(\mathbb{C})$ be the channel $\Phi_U(X) = UXU^*$ for $U \in U(2)$. Every Φ_U has Kraus rank one, and the set of distinct channels of this form is in bijection with $PU(2)$, which is not a point. The assignment from channels to Kraus rank is therefore not injective under equality of maps, and Kraus rank is not a classifying invariant for this family under that relation.*

Proof. For each U the single Kraus operator U realizes Φ_U , and no shorter dilation exists because Φ_U is nonzero. So the Kraus rank is one for every U . Two operators U, U' give the same map exactly when $UXU^* = U'XU'^*$ for all $X \in M_2(\mathbb{C})$, which forces $U' = \lambda U$ for a scalar λ of unit modulus. The distinct channels of this form are therefore parametrized by $PU(2)$, which is positive dimensional, and the fiber of the rank map over 1 is not a point. $\square$

The choice of equivalence in Theorem 12.2 is load-bearing and the conclusion changes with it. Under simultaneous unitary conjugacy the classes of these channels are the conjugacy classes of $PU(2)$, parametrized by a rotation angle in $[0, \pi]$, and rank is still not complete on them. Under two-sided conjugacy by automorphisms of $M_2(\mathbb{C})$ every unitary channel is isomorphic to the identity channel, so this particular family is not an isomorphism counterexample in that sense. What the proposition establishes is the narrow and sufficient claim that one scalar of a normal form does not classify: recording that scalar as though it were an isomorphism classification would produce a false row in the family table.

13 Family VIII: tensor categories and compositional structures

This family represents quantum information. Its objects are tensor categories, its morphisms are tensor functors, and its composition rule is the monoidal product.

Two islands are classified under explicit hypotheses. Symmetric tensor $*$ -categories with conjugates, direct sums, subobjects, and $\text{End}(1) = \mathbb{C}$ are equivalent to finite-dimensional unitary representations of a compact supergroup, unique up to isomorphism [12, 16]. Symmetric tensor categories of subexponential growth in characteristic zero are supergroup representation categories [13]. Both statements carry their hypotheses, and dropping a hypothesis produces a false statement.

The general case is open. Every fusion category is the representation category of a weak Hopf algebra, and the classification of fusion categories is open [14]. The constraint on any equivalence used in this family is a theorem and not a convention.

Theorem 13.1 (Structure preservation is load-bearing). *There exist finite groups that are not isomorphic and whose representation categories are equivalent as plain tensor categories. Consequently an equivalence that forgets the symmetric structure cannot recover the group, and any classification in this family must state which layer of structure its equivalence preserves [36].*

Theorem 13.1 is a negative theorem and it is used here as a discipline rule: an equivalence record that does not name the preserved structure is incomplete.

14 Family IX: discrete combinatorial structures

This family represents lattice and network models. Its state spaces are configuration spaces on a graph or lattice, its observables are local functions, and its composition rule is disjoint union or tensor product.

The proved results are regime-restricted and the impossibility results are not. Lattice gauge results hold in restricted regimes and do not by themselves close a continuum limit. The obstructions recorded against this family come from wild representation theory, and the underlying mathematics is commutative algebra and equivariant homotopy theory, which the taxonomy places in Family X. On the commutative side, rank-two isomorphism classes of maximal Cohen-Macaulay modules over the generic determinantal ring admit no parameterization by a finite-dimensional variety [32]. On the equivariant side, the compact modules over the equivariant Eilenberg-MacLane spectrum $H\mathbb{F}_p$ at an odd prime p are wild [34], and the qualifier is load-bearing: modules over the group ring $\mathbb{F}_p[C_p]$ have finite representation type, so the claim is false if the spectrum coefficient is dropped. A companion boundary is algorithmic and lives in this family: an isomorphism test running in polynomial time exists for groups that are tame extensions, so the tame and wild regimes are separated by an algorithmic boundary as well [33].

Corollary 14.1. *By Theorem 3.2, a sector whose classes admit no parameterization by a finite-dimensional variety admits no invariant-complete record whose range is a finite-dimensional variety. The obstruction is a statement about the range of any candidate invariant.*

14.1 A lattice gauge configuration

Take a finite graph Λ with a compact structure group G and a configuration space of G -valued edge assignments modulo the gauge action at the vertices. The record is: state space the quotient configuration space, observables the gauge-invariant functions generated by Wilson loops, dynamics the transfer operator, symmetries the gauge group, and composition the disjoint union of graphs.

The tame side of this sector admits named invariants: the holonomy of each closed loop, and the character of the representation assigned to it. The wild side does not, and Theorem 14.1 records what that absence means for any candidate invariant. Algebraic quantum field theory axiomatizes a theory as a local net of operator algebras [15], and that framework is where the lattice sector of this

family meets the operator-algebra family. Lattice results proved in restricted regimes do not by themselves close the continuum limit, and the invariant record of the limiting object is a separate matter, as Theorem 5.3 records.

15 Family X: derived and homotopical structures

This family represents higher gauge and derived field theories. Its objects are derived categories or homotopical enhancements, and its morphisms are derived equivalences.

Every finite-dimensional algebra over an algebraically closed field falls on one side of the tame-wild dichotomy for derived categories [31]. This family is the home of the wild-side obstructions that Family IX inherits: the non-parameterization of rank-two isomorphism classes of maximal Cohen-Macaulay modules over the generic determinantal ring [32], and the wildness of the compact modules over $H\mathbb{F}_p$ at an odd prime [34]. Beyond the dichotomy no classification is claimed.

Remark 15.1 (Dichotomy as content). A dichotomy theorem classifies the objects of a theory into two classes and classifies neither class internally. The row for this family in the chart below is complete at the level of the dichotomy and open at the level of the two sides.

16 Exceptional structures

The chart of families is not a chart of everything. Two kinds of additional structure sit beside it: genuinely exceptional objects, and distinguished objects that are canonical rather than exceptional. Both kinds are simple or indecomposable in their ambient categories, and that is the property the table below uses: the twenty-six sporadic simple groups together with the cyclic, alternating, and Lie-type families [1, 2, 3]; the exceptional Lie algebras G_2 , F_4 , E_6 , E_7 , E_8 [4]; the type III_0 and type III_1 factors and the hyperfinite II_1 factor [8]; the eight Thurston geometries [35] and the bounded list of four-dimensional model geometries that Filipkiewicz enumerated, recorded in the survey literature beside the unrecognizability results of [30]; and pairs of non-isomorphic finite groups with equivalent plain-tensor representation categories [36].

Table 2: Collections of simple or indecomposable objects recorded beside the families. The last column states why none of them is closed under direct sum; that property is about simplicity and not about exceptionality.

Collection	What it is	Why it is not closed under direct sum
Sporadic simple groups	Finite simple groups outside the cyclic, alternating, and Lie-type families [1, 2]	Direct sum or direct product of two such groups is not simple
Exceptional Lie algebras	The five simple Lie algebras of types G_2 , F_4 , E_6 , E_7 , E_8 [4]	Direct sum of two of them is not simple
Distinguished factors	Type III_0 , type III_1 , and the hyperfinite II_1 factor [8], none of them exceptional	Direct sum of any two factors has nontrivial center and is not a factor
Thurston geometries	The eight model geometries in dimension three [35]	Connected sum of geometrized manifolds need not carry a single geometry
Isocategorical pairs	Non-isomorphic finite groups with equivalent plain-tensor categories [36]	Recovering the group from the category fails without symmetric structure

Physical theories should not be assumed to possess a finite collection of exceptional objects analogous to the twenty-six sporadic groups. The table records a structural fact that holds for simple objects of any kind and is not specific to the exceptional ones: the direct sum of two simple objects is not simple, whether the two are exceptional or regular. The next proposition states that fact in the record setting.

Proposition 16.1 (Irreducible classes do not close). *Let $\mathcal{K}$ contain direct sum and let $\mathcal{E}$ be a nonempty class of nonzero objects each irreducible under direct sum. Then $\mathcal{E}$ is not closed under direct sum and $\langle \mathcal{E} \rangle_{\mathcal{K}}$ strictly contains $\mathcal{E}$. The conclusion is one of non-closure and not of non-generation: $\mathcal{E}$ generates $\langle \mathcal{E} \rangle_{\mathcal{K}}$ by definition of closure, so a generating class may consist of irreducibles alone.*

Proof. Take $A \in \mathcal{E}$ nonzero. The pair (A, A) witnesses a decomposition of $A \oplus A$ under the direct-sum construction with irreducible factors A and A , so $A \oplus A$ is reducible and is therefore not in $\mathcal{E}$. It is in $\langle \mathcal{E} \rangle_{\mathcal{K}}$ by closure under direct sum. The inclusion $\mathcal{E} \subsetneq \langle \mathcal{E} \rangle_{\mathcal{K}}$ is strict, which is the non-closure claim. The generation claim runs the other way and is immediate: $\langle \mathcal{E} \rangle_{\mathcal{K}}$ is by definition the smallest class containing $\mathcal{E}$ and closed under $\mathcal{K}$, so $\mathcal{E}$ generates it. $\square$

Two consequences follow and they are different. A class of irreducibles is never a fixed point of the closure operator whenever $\mathcal{K}$ contains direct sum, so taking a closure always adds objects. A class of irreducibles can nevertheless be a generating set, and the finite-dimensional matrix algebras are the example: each $M_n(\mathbb{C})$ is irreducible under direct sum and the list $\{M_n(\mathbb{C}) : n \geq 1\}$ generates every unital finite-dimensional complex C^* -algebra under direct sum, by Theorem 7.1.

17 Independence and reconstructability

A natural hope is that the ten families are logically independent, so that no family can be rebuilt from another. What the literature proves is a mixture of reconstruction theorems and counterexamples to reconstruction.

Table 3: Reconstruction and non-reconstruction results. The first four rows relate distinct families; the fifth compares two sectors of one family. Each row states the verdict and the structure the equivalence preserves.

Relation	Verdict	Preserved structure
Commutative unital C^* -algebras and compact Hausdorff spaces	Reconstructable	The Gelfand correspondence preserves the lattice of closed ideals [23, 24]
Boolean algebras and zero-dimensional compact spectra	Reconstructable	Stone duality preserves the order and the topology [25]
Symmetric tensor categories with conjugates and compact supergroups	Reconstructable	Doplicher-Roberts reconstruction preserves the symmetric structure [12, 16]
Plain tensor categories and the underlying finite group	Not reconstructable	No symmetric structure is preserved, and the group is lost [36]
Classical and quantum monotone metrics, both in Family V	Not identical	Uniqueness holds classically and fails in the quantum case [28]
Symplectic and operator-algebra families	Open	Deformation quantization is a program and not an equivalence

Definition 17.1 (Independence of families). Two structural families $\mathfrak{F}$ and $\mathfrak{G}$ are *logically independent* when neither is a special case of the other and neither is reconstructable from the other: there is no full embedding of one into the other, and no correspondence of the kind in Table 3 that determines one from the other. Independence in this sense is a statement about the families as classes of records and not about any particular system.

Proposition 17.2 (Invariants do not imply independence). *Let $\mathfrak{F}$ and $\mathfrak{G}$ be two structural families each admitting a classifying invariant. Then $\mathfrak{F}$ and $\mathfrak{G}$ need not be logically independent. In particular the commutative unital C^* -algebras inside the operator-algebra family and the compact Hausdorff spaces inside the geometric family each carry a classifying invariant, and neither is independent of the other.*

Proof. The Gelfand correspondence gives a bijection between commutative unital C^* -algebras and compact Hausdorff spaces that is an anti-equivalence of categories [24, 23]. The spectrum is a classifying invariant for the space side and the algebra is a classifying invariant for the algebra side, and each determines the other. Two families related by such a correspondence cannot be logically independent, because every statement about one transfers to the other. The existence of a classifying invariant in each family is therefore silent on the question of independence. $\square$

No theorem on record states that any two of the ten families are independent in the sense of Theorem 17.1. The honest chart entry for the independence question is open, and Theorem 17.2 explains why classifying invariants cannot settle it.

18 Systems in more than one family

A physical system is a record $(X, \mathcal{O}, \mathcal{D}, \mathcal{S}, \mathcal{C})$, and the families classify records by their mathematical structure. A single record can present structure from several families at once, and the taxonomy must account for that without pretending the families partition mathematics.

Proposition 18.1 (Mixed records are not defects). *There exist records whose state space is a symplectic manifold, whose observable algebra is a noncommutative C^* -algebra, and whose symmetry group is a Lie group. Such a record lies in the symplectic, operator-algebraic, and geometric families at once. Consequently the assignment of records to families is not a partition, and no argument may treat the ten families as mutually exclusive classes.*

Proof. Take M a compact Kähler manifold with integral Kähler form ω , let $L \rightarrow M$ be the corresponding positive line bundle, and let $\mathcal{H}_L$ be the space of holomorphic sections of L . For $f \in C^\infty(M)$ let T_f be the Toeplitz operator on $\mathcal{H}_L$ that multiplies by f and then projects back to $\mathcal{H}_L$. The C^* -algebra generated by the operators T_f is noncommutative whenever $\dim \mathcal{H}_L > 1$, since $T_f T_g - T_g T_f$ is nonzero for suitable f and g . The record whose state space is M with its symplectic form, whose observable algebra is this Toeplitz C^* -algebra, whose symmetries are the biholomorphisms of M lifting to L , and whose dynamics is the quantization of a Hamiltonian on M , presents symplectic, operator-algebraic, and geometric structure at once. If the families were mutually exclusive, this record would lie in none of them or in more than one, and the second option contradicts exclusivity. $\square$

The research goal at this level is to discover which distinctions are fundamental and which arise from different mathematical presentations. Theorem 18.1 is a small instance of that question, and it is the reason Theorem 2.4 defines a family by the structure it singles out rather than by membership of records.

Two further relations between families are worth naming. A family may be a special case of another: the commutative unital C^* -algebras are a full subcategory of the operator-algebra family and are equivalent to compact Hausdorff spaces, so the two presentations carry the same information [24]. A family may be generated from another under constructions, in the sense of Theorem 2.4. Both relations are structural statements about the families, and neither is a statement about the physical systems the families represent.

19 The taxonomy as an interface

Later volumes classify systems. The taxonomy is the place they record what they find, and the record format is fixed here so that results produced independently are comparable.

A classification record in a later volume has six fields: the family, the sector, the classifying invariant together with its range, the equivalence the invariant is complete for, the scope in which completeness holds, and the achievement kind of Theorem 3.6. A record with an empty invariant field is a record of an absence and not a classification.

Proposition 19.1 (Comparability of records). *Two classification records are comparable when they agree on family, sector, and equivalence. If they agree on those three fields and both name classifying invariants, then the two invariants are related by a bijection between their attained values on that sector.*

Proof. Both invariants are complete for the same equivalence on the same set of objects, so each induces a bijection between the set of classes and its own attained values. The composite of one bijection with the inverse of the other is a bijection between the two sets of attained values. $\square$

Theorem 19.1 is what makes the family rows usable across volumes. Two independently obtained invariants for the same sector are not competing claims; they are coordinate descriptions of the same class set, and each can be converted into the other once both are recorded. What is not automatic is agreement on scope, since two records may cover different sub-sectors. Scope comparison is therefore a manual step, and the family tables above state the scope of every row they carry.

The interface also fixes how an obstruction is recorded. An impossibility result names the sector, names the candidate invariant range it excludes, and cites the theorem that excludes it. An absence of record names the sector and states that no classification is on record. The two are different fields and never share a row, as Table 1 shows.

20 Scope of the family taxonomy

Four limits bound the use of the ten-family chart.

The families are candidates and not a completed classification. They are recorded as candidate organizing families, and some may be unified or shown to be special cases of others. The word “candidate” in Theorem 2.4 is load-bearing.

No completeness claim is available. Whether every admissible physical system lies in the closure of these ten families under the admissible constructions is unproved, and the conjectural identity that states it is a registered conjecture. Theorem 16.1 shows that a class of irreducibles is not closed under the constructions, so any completeness claim has to describe a closure and the objects it adds, not only a list of generators.

Status tokens are coarse. A single status per family hides the finite-dimensional, amenable, and wild sectors of the operator-algebra family at once, as Theorem 7.2 records. Table 5 gives one row

per family and states a scope limit in its last column; where the sectors of a family differ in kind, Tables 1 and 3 give the finer breakdown.

Independence is unrecorded. Theorem 17.2 shows the question is not settled by invariants, and no theorem settles it. Statements about the logical relations between families are therefore hypotheses in later volumes and not lemmas.

Table 4: Open questions the taxonomy raises, with the kind of result that would close each one.

Question	Present status	What would close it
Which families are fundamental rather than presentations of others	Open	A theorem reducing one family to another
Whether any two families are logically independent	Open	A non-derivability theorem in a stated sense
Whether the ten families generate all admissible systems	Conjecture	A proof of the closure identity, or an obstruction
A uniform parameterization of the finite simple groups	Open	An explicit uniform description of the classes
The global invariants of the symplectic family	Unrecorded	A named finite or parametric invariant list
A classification theorem for the semigroup family	None on record	A classification with a stated scope
Complete cases of the extended topological classification	Claimed	A proof closing the 2009 sketch in each case

21 Conclusion

The ten families are a working taxonomy. Each carries a classifying invariant or an explicit statement that none is on record, a set of classified islands with the scope they hold in, and an obstruction where an obstruction is proved. The invariant discipline of Section 3 is the method by which later volumes record their own classification results: one named invariant, one named range, one named equivalence, one named scope.

The four results set the terms of use. Classification shape follows invariant shape (Theorem 3.2), classification in the symplectic family is global (Theorem 3.4), a class of irreducibles is not closed under the admissible constructions (Theorem 16.1), and classifying invariants do not establish independence between families (Theorem 17.2). All four are proved here in full. The fourth is also a standing warning against reading the chart as a decomposition of mathematics.

What remains open is stated plainly. Which of the ten families are fundamental rather than presentations of others is not known. Whether any two families are logically independent is not known. Whether the ten families generate all admissible physical systems under the admissible constructions is a conjecture. The taxonomy is a place to record classification results, and it is not itself the classification.

A Computational evidence

The code accompanying this paper supplies executable models of the taxonomy and small computational checks of two claims. Computational evidence is evidence and never proof. The executable models are written against the family, invariant, and construction classes of the typed taxonomy

layer, and they are validated against the specifications recorded in the paper. They are not machine checked.

References

- [1] *Classification of finite simple groups*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Classification_of_finite_simple_groups.
- [2] R. Solomon, *A brief history of the classification of the finite simple groups*, Bulletin of the American Mathematical Society **38** (2001) 315-352. Available at <https://www.ams.org/journals/bull/2001-38-03/S0273-0979-01-00909-0/>.
- [3] *ATLAS of Finite Groups*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/ATLAS_of_Finite_Groups.
- [4] *Simple Lie group*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Simple_Lie_group.
- [5] *Theorem of the highest weight*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Theorem_of_the_highest_weight.
- [6] *Darboux's theorem*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Darboux%27s_theorem.
- [7] *Poisson manifold*. WikiJournal of Science **7** (2024) 6. Available at https://en.wikipedia.org/wiki/Poisson_manifold.
- [8] *Von Neumann algebra*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Von_Neumann_algebra.
- [9] *C^* -algebra*. Encyclopedia article. Available at [https://en.wikipedia.org/wiki/\\$C^*\\$-algebra](https://en.wikipedia.org/wiki/C^*-algebra).
- [10] *Stinespring dilation theorem*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Stinespring_dilation_theorem.
- [11] *Quantum channel*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Quantum_channel.
- [12] *Doplicher-Roberts reconstruction theorem*. nLab entry. Available at <https://ncatlab.org/nlab/show/Doplicher-Roberts+reconstruction+theorem>.
- [13] *Deligne's theorem on tensor categories*. nLab entry. Available at <https://ncatlab.org/nlab/show/Deligne%27s+theorem+on+tensor+categories>.
- [14] *Fusion category*. nLab entry. Available at <https://ncatlab.org/nlab/show/fusion+category>.
- [15] *Haag-Kastler axioms*. nLab entry. Available at <https://ncatlab.org/nlab/show/Haag-Kastler+axioms>.
- [16] *DHR superselection theory*. nLab entry. Available at <https://ncatlab.org/nlab/show/DHR+superselection+theory>.

- [17] *Classifying space*. nLab entry. Available at <https://ncatlab.org/nlab/show/classifying+space>.
- [18] *Characteristic class*. nLab entry. Available at <https://ncatlab.org/nlab/show/characteristic+class>.
- [19] *Quantum anomaly*. nLab entry. Available at <https://ncatlab.org/nlab/show/quantum+anomaly>.
- [20] J. C. Baez and J. Dolan, *Higher-dimensional algebra and topological quantum field theory*, arXiv:q-alg/9503002 (1995). Available at <https://arxiv.org/abs/q-alg/9503002>.
- [21] J. Lurie, *On the classification of topological field theories*, arXiv:0905.0465 (2009). Available at <https://arxiv.org/abs/0905.0465>.
- [22] *Topological quantum field theory*. nLab entry. Available at <https://ncatlab.org/nlab/show/topological+quantum+field+theory>.
- [23] *Gelfand-Naimark theorem*. nLab entry. Available at <https://ncatlab.org/nlab/show/Gelfand-Naimark+theorem>.
- [24] *Gelfand duality*. nLab entry. Available at <https://ncatlab.org/nlab/show/Gelfand+duality>.
- [25] *Stone duality*. nLab entry. Available at <https://ncatlab.org/nlab/show/Stone+duality>.
- [26] E. Minguzzi and M. Sánchez, *The causal hierarchy of spacetimes*, arXiv:gr-qc/0609119 (2006). Available at <https://arxiv.org/abs/gr-qc/0609119>.
- [27] A. Galaev, *Holonomy groups of Lorentzian manifolds*, arXiv:1611.01551 (2016). Available at <https://arxiv.org/abs/1611.01551>.
- [28] D. Petz and C. Sudár, *Extending the Fisher metric to density matrices*, arXiv:quant-ph/0102132 (2001). Available at <https://arxiv.org/abs/quant-ph/0102132>.
- [29] *Adyan-Rabin theorem*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Adian%E2%80%93Rabin_theorem.
- [30] M. Tancer, *Simpler algorithmically unrecognizable 4-manifolds*, arXiv:2310.07421 (2023). Available at <https://arxiv.org/abs/2310.07421>.
- [31] V. Bekkert and Y. Drozd, *Tame-wild dichotomy for derived categories*, arXiv:math/0310352 (2003). Available at <https://arxiv.org/abs/math/0310352>.
- [32] R.-O. Buchweitz and G. J. Leuschke, *Factoring the adjoint and maximal Cohen-Macaulay modules over the generic determinant*, arXiv:math/0505315 (2005). Available at <https://arxiv.org/abs/math/0505315>.
- [33] J. A. Grochow and Y. Qiao, *Polynomial-time isomorphism test of groups that are tame extensions*, arXiv:1507.01917 (2015). Available at <https://arxiv.org/abs/1507.01917>.
- [34] J. F. Grevstad and C. May, *Equivariant $H\mathbb{F}_p$ -modules are wild*, arXiv:2509.17604 (2025). Available at <https://arxiv.org/abs/2509.17604>.

- [35] *Geometrization conjecture*. Encyclopedia article. Available at https://en.wikipedia.org/wiki/Geometrization_conjecture.
- [36] P. Etingof and S. Gelaki, *Isocategorical groups*, International Mathematics Research Notices **2001** (2001) 59-120. DOI 10.1155/s1073792801000046. Available at <https://doi.org/10.1155/s1073792801000046>.
- [37] D. McDuff, *Uncountably many II_1 factors*, Annals of Mathematics **90** (1969) 351-370. DOI 10.2307/1970730. Available at <https://doi.org/10.2307/1970730>.
- [38] R. Sasyk and A. Törnquist, *Borel reducibility and classification of von Neumann algebras*, arXiv:0909.2657 (2009). Available at <https://arxiv.org/abs/0909.2657>.
- [39] A. Weinstein, *The local structure of Poisson manifolds*, Journal of Differential Geometry **18** (1983) 523-557. DOI 10.4310/jdg/1214437787. Available at <https://doi.org/10.4310/jdg/1214437787>.
- [40] W. F. Stinespring, *Positive functions on C^* -algebras*, Proceedings of the American Mathematical Society **6** (1955) 211-216. DOI 10.1090/s0002-9939-1955-0069403-4. Available at <https://doi.org/10.1090/s0002-9939-1955-0069403-4>.

Table 5: The ten candidate structural families with their representative systems, classifying invariants, and achievement kind. The kind column uses the vocabulary of Theorem 3.6; where the sectors of a family differ, the row says so and Table 1 gives the breakdown.

Family	Defining data	Classifying invariant	Achievement kind	Scope limit
I	Symplectic or Poisson state space	Global symplectic topology	Mixed: parameterized locally, no record globally	No local invariants exist
II	Noncommutative observable algebra	Dimension vector, type and Connes data	Mixed: parameterized, and impossible in part	Finite-dimensional and amenable sectors only
III	Smooth manifold with geometry	Holonomy, causal class, geometry type	Mixed: parameterized in restricted regimes	Dimension four is unrecognizable
IV	Principal bundle with connection	Map to the classifying space	Complete and parameterized	Base and group must be numerable
V	Space of probability measures	Fisher metric up to constant	No record for the systems	Uniqueness is of the metric, not of the record
VI	Cobordism datum	Fully dualizable object	No record, with a claim on file	The 2009 proof is a sketch
VII	Operator semigroup or generator	None on record	No record	The Kraus form is a normal form, and rank is one scalar of it
VIII	Tensor category	Supergroup data with symmetric layer	Mixed: parameterized islands, no record generally	Hypotheses travel with the theorem
IX	Configuration space on a lattice	Tame or wild class	Mixed: impossible in wild parts	No finite-variety parameterization
X	Derived or homotopical category	Derived-tame or derived-wild class	No record beyond the dichotomy	Nothing is claimed beyond it

Table 6: Code evidence. Each row names a claim, the file that evidences it, and what the file demonstrates.

Row	Claim	Evidence file	What the file demonstrates
1	Kraus rank is constant on a positive-dimensional family of channels	<code>phased_kraus_rank.py</code>	Constructs unitary channels on M_2 and reports rank one for each
2	Dimension vectors enumerate the finite-dimensional sector at fixed dimension	<code>phased_dimension_vectors.py</code>	Enumerates dimension vectors of total dimension n and counts them
3	Each of the ten families has a declared instance of the taxonomy classes	<code>Instances.hs</code>	Declares one instance per family for the systems named in this paper
4	Family tags carry ontology identifiers	<code>Instances.hs</code>	Reports the identifier of each family tag through the taxonomy accessor