

Presentation Theory VI: Observable Quotients, Separation, and Reconstruction

Luca Blachi

June 2026

Abstract

This paper develops the observable-reconstruction layer of Presentation Theory. Earlier parts introduced presentation systems, controlled transfer, normal-form compilation, verifiable access, and fibre geometry. Here the primary objects are observables themselves. A family of observables produces numerical or algebraic data and, at the same time, induces a quotient of the object class by the relation of observational indistinguishability. A mathematical target also induces a quotient, by identifying objects with the same target value. Reconstruction is the comparison of these two quotients.

The formal core is a calculus of observable quotients. A profile F reconstructs a target τ exactly when its indistinguishability relation E_F is contained in the target kernel K_τ . Separation power is refinement of these relations. Post-processing gives an observational closure operator, and observables and equivalence relations form an order-reversing Galois correspondence. In effective filtered systems, finite-level observable closure gives computable outer approximations to true saturation; complete effective observable separation is equivalent to decidability of the corresponding bounded-presentation or equivalence problem. If a profile does not reconstruct a target, the quotient $\mathcal{C}/(K_\tau \vee E_F)$ is the most informative factor of the target still determined by the profile. Residual target ambiguity can then be measured by cardinality, entropy, or metric diameter. Relative reconstruction is governed by $E_F \cap K_\sigma \subseteq K_\tau$, so additional observables are best understood as resolving residual target ambiguity inside a base quotient.

The quantitative part records standard access models. Window budgets give a pairwise separation formula for the reconstruction threshold. Multi-resource budgets give upward-closed reconstruction regions and Pareto frontiers. Finite fingerprints reduce exactly to weighted set cover on target-distinct pairs, and the dual linear program gives packing lower bounds. Adaptive reconstruction is decision-tree reconstruction with target-homogeneous leaves. Structured residual regimes recover familiar algebra: linear residuals require spanning the dual, torsorial residuals require separating characters, algebraic residuals require the target field to lie in the field generated by the observables, logical residuals require definability in the declared fragment, and probabilistic observables are ordered by stochastic post-processing. Stability is measured by a reconstruction modulus; exact reconstruction need not be stable.

The final sections explain how this observable calculus interfaces with the other parts of the theory and with recurring applications: rank invariants in multiparameter persistence, principal Hecke data and class-group twist ambiguity, and finite hom-count or Weisfeiler–Leman completions for graph classes. These examples isolate the reusable quotient, residual, cost, and verification structure shared by many separating-family arguments.

1 Purpose

Presentation Theory studies mathematical access through declared systems of description, observation, verification, and transfer. A presentation describes an object. A realization map forgets the description and remembers the object. A normal form chooses representatives. A

transfer package moves access between contexts. Fibre geometry studies the internal structure left inside a realization fibre.

This paper studies the complementary external question:

which aspects of an object are determined by the observables we are allowed to read?

The answer is naturally quotient-theoretic. A family of observables separates some pairs of objects and fails to separate others. Hence it induces an equivalence relation of indistinguishability. A target, such as isomorphism type, semisimplification, a Hecke eigensystem, a rank invariant, or the truth value of a property, also induces an equivalence relation. The observable family reconstructs the target precisely when its indistinguishability relation is finer than the target relation.

The guiding formula is

$$F \text{ reconstructs } \tau \iff E_F \subseteq K_\tau.$$

Everything else in this paper refines this inclusion: how it behaves under post-processing, what part of the target remains determined when it fails, how much it costs to make it true, how finite reconstruction becomes set cover, how adaptive reconstruction becomes a decision tree, how residual linear or torsorial ambiguity is resolved, how stability is measured, and how the inclusion is transported or verified.

Relation with the preceding parts

Part I introduced observables as part of the basic language. Part II studied controlled transfer of observables and budgets. Part III studied normal forms. Part IV studied auditable verification under bounded access. Part V studied the geometry internal to realization fibres. The present paper makes observables the main object. It studies the quotient of the object class seen by a profile of observables and compares it with the quotient required by a target.

In short:

Part	Main role
I	presentation systems, costs, fibres, observables
II	controlled transfer between access contexts
III	normal-form compilation and canonical access
IV	verifiable access and audit lower bounds
V	internal geometry of realization fibres
VI	observable quotients, separation, and reconstruction

Notation and conventions

The letter $\mathcal{C}$ denotes the object class under study. A family F of observables is always interpreted relative to a declared observable system (Ω, χ, B) . The relation E_F is the equivalence relation of having the same F -profile. A target $\tau : \mathcal{C} \rightarrow T$ induces $K_\tau = \ker \tau$. If E and E' are equivalence relations on $\mathcal{C}$, then

$$E \vee E'$$

denotes their join in the lattice of equivalence relations: the smallest equivalence relation containing both. Equivalently, it is the transitive closure of $E \cup E'$. Quotients such as $\mathcal{C}/E$ are used in the ordinary set-theoretic sense when $\mathcal{C}$ is a set; in large or structured contexts they should be read as the corresponding quotient object in the declared universe or category.

The order convention is the following. A smaller equivalence relation carries more information, because fewer pairs are identified. Thus $E_F \subseteq E_G$ means that F is at least as separating as G . A factor of a target is a coarser target: σ is a factor of τ precisely when $K_\tau \subseteq K_\sigma$.

2 Observable systems

Definition 2.1 (Observable system). Let $\mathcal{C}$ be a class of objects. An observable system on $\mathcal{C}$ consists of:

$$(\Omega, \chi, B),$$

where Ω is a class of observables $\omega : \mathcal{C} \rightarrow A_\omega$, B is an ordered budget scale, and

$$\chi : \Omega \rightarrow B$$

assigns an observation cost to each observable.

The budget scale may be $\mathbb{N}$, $\mathbb{R}_{\geq 0}$, a product such as $\mathbb{N}^r$, or another ordered resource set. Typical coordinates measure degree, arity, precision, sample count, operator norm, test-object size, query depth, memory, or verification weight.

Example 2.2 (Standard observables). Examples include:

Context	Object	Observable
graphs	G	$\# \text{Hom}(T, G)$
persistence	M	$\text{rank}(M(a) \rightarrow M(b))$
representations	ρ	$\text{Tr}(\rho(g))$
groups	G	$\Pr_{x,y \in G}([x, y] = 1)$
Hecke systems	f	$a_a(f)$
finite structures	M	$1_{M \models \varphi}$
algebraic geometry	$x \in X$	$f(x)$

The codomain A_ω may be numerical, algebraic, logical, categorical, probabilistic, or topological.

Definition 2.3 (Profile). For a family $F \subseteq \Omega$, the observational profile of F is the map

$$\Phi_F : \mathcal{C} \rightarrow \prod_{\omega \in F} A_\omega, \quad \Phi_F(x) = (\omega(x))_{\omega \in F}.$$

If $F = \{\omega_1, \dots, \omega_m\}$, then

$$\Phi_F(x) = (\omega_1(x), \dots, \omega_m(x)).$$

Definition 2.4 (Budget filtration). For $b \in B$, define the available observables at budget b by

$$\Omega_{\leq b} = \{\omega \in \Omega : \chi(\omega) \leq b\}.$$

The corresponding profile and indistinguishability relation will be denoted by

$$\Phi_b = \Phi_{\Omega_{\leq b}}, \quad E_b = E_{\Omega_{\leq b}}.$$

When $b \leq b'$, one has $\Omega_{\leq b} \subseteq \Omega_{\leq b'}$, so the equivalence relations move in the opposite direction:

$$E_{b'} \subseteq E_b.$$

Thus increasing the observation budget refines the observational quotient.

3 Effective Observable Closure and Separation

The quotient language above becomes especially sharp for filtered effective presentation systems. Let

$$\Gamma = (D, \kappa, (E_t))$$

be a filtered effective presentation system as in Part I, with equivalence relation $E = \bigcup_t E_t$. An effective observable family on D is a family

$$\mathcal{O} = (\Lambda, \lambda, (O_\alpha)_{\alpha \in \Lambda})$$

such that $\Lambda \subseteq \Sigma^*$ is decidable, $\lambda : \Lambda \rightarrow \mathbb{N}$ is computable with finite computable sublevel sets

$$\Lambda_{\leq s} = \{\alpha : \lambda(\alpha) \leq s\},$$

each $O_\alpha : D \rightarrow A_\alpha$ is computable uniformly in α , equality in each A_α is decidable, and

$$d E e \implies O_\alpha(d) = O_\alpha(e)$$

for every α .

For finite $S \subseteq D$, define the finite-level observable closure

$$\text{Cl}_{\mathcal{O},s}(S) = \{d \in D : \exists e \in S \text{ such that } O_\alpha(d) = O_\alpha(e) \text{ for all } \alpha \in \Lambda_{\leq s}\}.$$

The full observable closure is

$$\text{Cl}_{\mathcal{O},\infty}(S) = \bigcap_s \text{Cl}_{\mathcal{O},s}(S).$$

The E -saturation of S is

$$\text{Sat}_E(S) = \{d \in D : \exists e \in S, d E e\}.$$

For every s ,

$$\text{Sat}_E(S) \subseteq \text{Cl}_{\mathcal{O},s}(S),$$

because observables are constant on E -classes.

Proposition 3.1 (Full closure is observable saturation). *Let $S \subseteq D$ be finite. Define $d \equiv_{\mathcal{O}} e$ if all observables in $\mathcal{O}$ have the same value on d and e . Then*

$$\text{Cl}_{\mathcal{O},\infty}(S) = \{d \in D : \exists e \in S, d \equiv_{\mathcal{O}} e\}.$$

Proof. The inclusion from right to left is immediate. Conversely, suppose $d \in \text{Cl}_{\mathcal{O},\infty}(S)$. For every s , choose $e_s \in S$ with the same observables as d up to level s . Since S is finite, some $e \in S$ occurs for infinitely many, hence unboundedly many, values of s . Given any observable O_α , choose such an s with $\lambda(\alpha) \leq s$. Then $O_\alpha(d) = O_\alpha(e)$. Thus $d \equiv_{\mathcal{O}} e$. $\square$

3.1 Bounded-presentation separation

For a cost bound b , the low-cost description set $D_{\leq b}$ is finite and computable. Since

$$(d, b) \in P_\Gamma \iff d \in \text{Sat}_E(D_{\leq b}),$$

finite observable closure gives a sound lower-bound test:

$$d \notin \text{Cl}_{\mathcal{O},s}(D_{\leq b}) \implies C_\Gamma(d) > b.$$

Define the residual set

$$R_{\Gamma,\mathcal{O}}(n, b, s) = \{d \in D : |d| \leq n, d \in \text{Cl}_{\mathcal{O},s}(D_{\leq b}), C_\Gamma(d) > b\}.$$

An observable family is **presentation-complete** for Γ if, for every d and b with $C_\Gamma(d) > b$, there is some finite level s such that

$$d \notin \text{Cl}_{\mathcal{O},s}(D_{\leq b}).$$

When this holds, define the separation profile

$$S_{\Gamma,\mathcal{O}}(n,b) = \min\{s : R_{\Gamma,\mathcal{O}}(n,b,s) = \emptyset\}.$$

Theorem 3.2 (Effective observable separation and bounded presentation). *For a filtered effective presentation system Γ , the following are equivalent:*

- (i) P_Γ is decidable;
- (ii) there exists an effective observable family that is presentation-complete.

Moreover, if a fixed effective observable family is presentation-complete, then $S_{\Gamma,\mathcal{O}}$ is computable.

Proof. Assume first that $\mathcal{O}$ is presentation-complete. The positive instances of P_Γ are c.e. by the finite-stage equivalence filtration: enumerate t and search for $e \in D_{\leq b}$ with $d E_t e$. The negative instances are also c.e.: enumerate s and test the decidable condition

$$d \notin \text{Cl}_{\mathcal{O},s}(D_{\leq b}).$$

Soundness gives correctness, and presentation-completeness gives termination on every negative instance. Thus P_Γ is decidable.

Conversely, suppose P_Γ is decidable. For each $m \in \mathbb{N}$, define

$$J_m(d) = \begin{cases} 1, & C_\Gamma(d) \leq m, \\ 0, & C_\Gamma(d) > m. \end{cases}$$

Since P_Γ is decidable, the functions J_m are computable uniformly in m , and they are constant on E -classes. If $C_\Gamma(d) > b$, then $J_b(d) = 0$, while every $e \in D_{\leq b}$ satisfies $J_b(e) = 1$. Hence the family (J_m) separates every negative bounded-presentation instance at finite level.

Finally, if $\mathcal{O}$ is a fixed effective presentation-complete family, then decidability of P_Γ allows one to decide, for fixed (n,b,s) , whether $R_{\Gamma,\mathcal{O}}(n,b,s)$ is empty by checking the finite ball $|d| \leq n$. Searching for the least s with empty residual set computes $S_{\Gamma,\mathcal{O}}(n,b)$. $\square$

3.2 Equivalence separation

An effective observable family is **equivalence-complete** if

$$d \equiv_{\mathcal{O}} e \iff d E e.$$

For non-equivalent pairs, define the pair separation level

$$\tau_{\Gamma,\mathcal{O}}(d,e) = \min\{s : \exists \alpha \in \Lambda_{\leq s}, O_\alpha(d) \neq O_\alpha(e)\},$$

when it exists, and define

$$T_{\Gamma,\mathcal{O}}(N) = \max\{\tau_{\Gamma,\mathcal{O}}(d,e) : |d| + |e| \leq N, \neg(d E e)\}.$$

Theorem 3.3 (Effective observable classification and equivalence). *For a filtered effective presentation system Γ , the following are equivalent:*

- (i) EQ_Γ is decidable;

(ii) there exists an effective observable family that is equivalence-complete.

If a fixed effective observable family is equivalence-complete, then $T_{\Gamma, \mathcal{O}}$ is computable.

Proof. If $\mathcal{O}$ is equivalence-complete, then EQ_{Γ} is c.e. from the filtration E_t , and its complement is c.e. by searching for a finite-level observable separating d and e . Thus EQ_{Γ} is decidable.

Conversely, if EQ_{Γ} is decidable, define for each $a \in D$

$$J_a(d) = \begin{cases} 1, & d E a, \\ 0, & \neg(d E a). \end{cases}$$

These observables are computable uniformly in a , constant on E -classes, and separate any non-equivalent pair d, e , for instance by J_d . This gives an effective equivalence-complete family after assigning a computable complexity such as $\lambda(J_a) = |a|$.

For a fixed equivalence-complete family, decidability of EQ_{Γ} lets one enumerate the finitely many non-equivalent pairs with $|d| + |e| \leq N$, find their least separating levels, and take the maximum. Hence $T_{\Gamma, \mathcal{O}}$ is computable. $\square$

Corollary 3.4 (Completeness is already a decision procedure). *If P_{Γ} is undecidable, no effective observable family is presentation-complete. If EQ_{Γ} is undecidable, no effective observable family is equivalence-complete.*

3.3 Finite basis versus correctness

Let $\mathcal{S}$ be a declared class of finite or effectively represented subsets of D . Finite-basis behavior and correctness are distinct conditions.

Finite basis asks whether, for $S \in \mathcal{S}$,

$$\text{Cl}_{\mathcal{O}, \infty}(S) = \text{Cl}_{\mathcal{O}, s}(S)$$

for some finite s , preferably with s bounded by a declared resource function of the parameters of S .

Correctness asks whether the observable quotient is fine enough for the intended equivalence:

$$\text{Cl}_{\mathcal{O}, \infty}(S) = \text{Sat}_E(S).$$

The first condition is a noetherian or finite-generation statement about the observable family. The second is a reconstruction statement about the target equivalence. Their conjunction gives finite observable separation of S .

For a chosen growth class $\mathcal{F}$, this gives the trichotomy:

- (i) quotient failure: correctness fails for some $S \in \mathcal{S}$;
- (ii) $\mathcal{F}$ -bounded finite basis: correctness holds and the required finite level is bounded in $\mathcal{F}$;
- (iii) unbounded finite basis: correctness holds, but no bound in $\mathcal{F}$ controls the required level.

At the level of all computable functions and finite windows, the third case disappears for effective complete observable families. It becomes meaningful for smaller growth classes, non-effective observable systems, or infinite definable families where finite-window search is no longer automatic.

4 Observable quotients and targets

Definition 4.1 (Indistinguishability relation). Let $F \subseteq \Omega$. The indistinguishability relation induced by F is

$$xE_Fy \iff \Phi_F(x) = \Phi_F(y).$$

Equivalently,

$$xE_Fy \iff \omega(x) = \omega(y) \text{ for every } \omega \in F.$$

Thus

$$E_F = \ker \Phi_F.$$

The quotient

$$q_F : \mathcal{C} \rightarrow \mathcal{C}/E_F$$

is the observable quotient induced by F .

Remark 4.2. The concrete profile $\Phi_F(x)$ may contain more coding information than the quotient $\mathcal{C}/E_F$. For reconstruction, only the induced separation relation matters. The quotient records exactly which objects have been separated and which have not.

Definition 4.3 (Target). A target is a map

$$\tau : \mathcal{C} \rightarrow T.$$

It induces the target kernel

$$K_\tau = \{(x, y) \in \mathcal{C} \times \mathcal{C} : \tau(x) = \tau(y)\}.$$

Equivalently, $K_\tau = \ker \tau$.

A target may be a complete invariant, a partial invariant, a property, a quotient, a semisimplification, a normal-form value, an orbit map, or a finite decision problem. The target relation K_τ identifies objects that are equivalent for the question being asked.

Example 4.4. If $\tau(G)$ is the isomorphism class of a finite graph G , then K_τ is graph isomorphism. If $\tau(\rho) = \rho^{ss}$, then K_τ identifies representations with the same semisimplification. If $\tau(M)$ is the rank invariant of a persistence module, then K_τ identifies modules with the same rank profile.

5 Universal reconstruction

Definition 5.1 (Reconstruction). A family $F \subseteq \Omega$ reconstructs a target $\tau : \mathcal{C} \rightarrow T$ if the value of $\tau(x)$ is determined by the profile $\Phi_F(x)$. Equivalently, for all $x, y \in \mathcal{C}$,

$$\Phi_F(x) = \Phi_F(y) \implies \tau(x) = \tau(y).$$

Theorem 5.2 (Universal reconstruction theorem). Let $F \subseteq \Omega$, and let $\tau : \mathcal{C} \rightarrow T$ be a target. The following are equivalent:

- (a) F reconstructs τ .
- (b) $E_F \subseteq K_\tau$.
- (c) τ is constant on the fibres of Φ_F .
- (d) There exists a map $R : \Phi_F(\mathcal{C}) \rightarrow T$ such that

$$\tau = R \circ \Phi_F.$$

(e) There exists a map $\bar{\tau} : \mathcal{C}/E_F \rightarrow T$ such that

$$\tau = \bar{\tau} \circ q_F.$$

Proof. By definition, F reconstructs τ exactly when

$$\Phi_F(x) = \Phi_F(y) \Rightarrow \tau(x) = \tau(y)$$

for all x, y . The hypothesis $\Phi_F(x) = \Phi_F(y)$ is exactly $x E_F y$, and the conclusion $\tau(x) = \tau(y)$ is exactly $x K_\tau y$. Hence reconstruction is equivalent to $E_F \subseteq K_\tau$.

The inclusion $E_F \subseteq K_\tau$ says precisely that τ is constant on every fibre of Φ_F , proving the equivalence with (c).

If τ is constant on the fibres of Φ_F , define

$$R(\Phi_F(x)) = \tau(x).$$

This is well-defined because any two choices of x with the same profile have the same target value. Then $\tau = R \circ \Phi_F$, giving (d). Conversely, if such an R exists, equal F -profiles imply equal target values.

The same argument applies to the quotient map q_F . The map

$$\bar{\tau}([x]_{E_F}) = \tau(x)$$

is well-defined exactly when τ is constant on E_F -classes. This gives (e), and the converse is immediate. $\square$

Remark 5.3. Reconstruction and normal forms answer different questions. If τ is an isomorphism class, reconstructing τ determines the class. Choosing a representative requires a section of the quotient and belongs to the normal-form layer.

Remark 5.4. Reconstruction and verifiable access are adjacent layers. The inclusion $E_F \subseteq K_\tau$ is the mathematical reconstruction statement. A bounded checker may additionally require proof data showing that the observed profile cell is contained in a target fibre. This is the interface with audit systems.

6 Separation power and observational closure

Definition 6.1 (Separation preorder). For $F, G \subseteq \Omega$, write

$$F \succeq G$$

if F separates at least all pairs separated by G . Formally,

$$F \succeq G \iff E_F \subseteq E_G.$$

The direction is contravariant: a smaller indistinguishability relation means greater separating power.

Definition 6.2 (Reconstructible targets). Let

$$\text{Rec}(F) = \{\tau : \mathcal{C} \rightarrow T \text{ for some } T : E_F \subseteq K_\tau\}$$

be the class of targets reconstructed by F .

Theorem 6.3 (Dominance theorem). For $F, G \subseteq \Omega$, the following are equivalent:

- (a) $F \succeq G$.
- (b) $E_F \subseteq E_G$.
- (c) The quotient map $q_G : \mathcal{C} \rightarrow \mathcal{C}/E_G$ factors through $q_F : \mathcal{C} \rightarrow \mathcal{C}/E_F$.
- (d) $\text{Rec}(G) \subseteq \text{Rec}(F)$.

Proof. The equivalence of (a) and (b) is the definition.

If $E_F \subseteq E_G$, every E_F -class is contained in an E_G -class. Therefore q_G is constant on the fibres of q_F , so there is a unique map

$$h : \mathcal{C}/E_F \rightarrow \mathcal{C}/E_G$$

with $q_G = h \circ q_F$. Conversely, if q_G factors through q_F , then equal q_F -values imply equal q_G -values, so $E_F \subseteq E_G$.

Assume $E_F \subseteq E_G$. If $\tau \in \text{Rec}(G)$, then $E_G \subseteq K_\tau$. Hence $E_F \subseteq K_\tau$, so $\tau \in \text{Rec}(F)$. Thus $\text{Rec}(G) \subseteq \text{Rec}(F)$.

Conversely, assume $\text{Rec}(G) \subseteq \text{Rec}(F)$. The quotient target $q_G : \mathcal{C} \rightarrow \mathcal{C}/E_G$ is reconstructed by G . By hypothesis it is reconstructed by F . Therefore $E_F \subseteq K_{q_G} = E_G$. $\square$

Definition 6.4 (Observational closure). For $F \subseteq \Omega$, define

$$\overline{F} = \{\eta \in \Omega : E_F \subseteq K_\eta\}.$$

Thus $\overline{F}$ is the family of allowed observables whose values are determined by the profile F .

Proposition 6.5 (Closure properties). *The assignment $F \mapsto \overline{F}$ satisfies:*

- (a) $F \subseteq \overline{F}$.
- (b) $F \subseteq G \Rightarrow \overline{F} \subseteq \overline{G}$.
- (c) $\overline{\overline{F}} = \overline{F}$.
- (d) $E_{\overline{F}} = E_F$.

Proof. If $\omega \in F$, then equal F -profiles imply equal ω -values, so $E_F \subseteq K_\omega$. Hence $F \subseteq \overline{F}$.

If $F \subseteq G$, then $E_G \subseteq E_F$. Thus any observable constant on E_F -classes is also constant on E_G -classes, so $\overline{F} \subseteq \overline{G}$.

Since $F \subseteq \overline{F}$, one has $E_{\overline{F}} \subseteq E_F$. On the other hand, every observable in $\overline{F}$ is constant on E_F -classes, so $E_F \subseteq E_{\overline{F}}$. Hence $E_{\overline{F}} = E_F$.

Finally,

$$\overline{\overline{F}} = \{\eta : E_{\overline{F}} \subseteq K_\eta\} = \{\eta : E_F \subseteq K_\eta\} = \overline{F}.$$

$\square$

Theorem 6.6 (Observable-equivalence Galois correspondence). *Let E be an equivalence relation on $\mathcal{C}$. Define*

$$\text{Obs}(E) = \{\omega \in \Omega : E \subseteq K_\omega\}.$$

Then for every $F \subseteq \Omega$,

$$F \subseteq \text{Obs}(E) \iff E \subseteq E_F.$$

Proof. If $F \subseteq \text{Obs}(E)$, then each $\omega \in F$ is constant on E -classes. Hence xEy implies $\omega(x) = \omega(y)$ for all $\omega \in F$, so $x E_F y$. Thus $E \subseteq E_F$.

Conversely, if $E \subseteq E_F$ and $\omega \in F$, then $x E y$ implies $x E_F y$, hence $\omega(x) = \omega(y)$. Therefore $\omega \in \text{Obs}(E)$, and $F \subseteq \text{Obs}(E)$. $\square$

7 Incomplete profiles and determined target factors

A profile may be incomplete for the intended target while still determining a meaningful factor of that target. The useful question is then to identify that factor.

Definition 7.1 (Determined target factor). Let $F \subseteq \Omega$, and let $\tau : \mathcal{C} \rightarrow T$ be a target. The target factor determined by F is the quotient

$$Q_F(\tau) = \mathcal{C}/(K_\tau \vee E_F),$$

where $K_\tau \vee E_F$ is the join of equivalence relations, namely the smallest equivalence relation containing both K_τ and E_F .

Theorem 7.2 (Universal property of the determined factor). *The quotient $Q_F(\tau)$ is the most informative factor of τ that is reconstructed by F . More precisely, let $\sigma : \mathcal{C} \rightarrow S$ be any target such that:*

- (a) σ is a factor of τ , equivalently $K_\tau \subseteq K_\sigma$;
- (b) σ is reconstructed by F , equivalently $E_F \subseteq K_\sigma$.

Then σ factors through the quotient map

$$\mathcal{C} \rightarrow Q_F(\tau).$$

Proof. Let

$$q : \mathcal{C} \rightarrow \mathcal{C}/(K_\tau \vee E_F)$$

be the quotient map. Since $K_\tau \subseteq K_\tau \vee E_F$, the map q is a factor of τ . Since $E_F \subseteq K_\tau \vee E_F$, the target q is reconstructed by F .

Now let σ satisfy the two stated conditions. Then K_σ contains both K_τ and E_F . Since $K_\tau \vee E_F$ is the smallest equivalence relation containing them, one has

$$K_\tau \vee E_F \subseteq K_\sigma.$$

Therefore σ is constant on the fibres of q , so σ factors uniquely through q . $\square$

Remark 7.3. The order language is delicate. The quotient by $K_\tau \vee E_F$ is coarser than the original target quotient $\mathcal{C}/K_\tau$, because the profile may force additional identifications. Among all factors of τ still determined by F , it is the finest one.

Definition 7.4 (Residual target set). For $x \in \mathcal{C}$, define the residual target set left by F at x by

$$R_F^\tau(x) = \{\tau(y) : y \in \mathcal{C}, \Phi_F(y) = \Phi_F(x)\}.$$

Equivalently,

$$R_F^\tau(x) = \tau(E_F[x]).$$

Proposition 7.5 (Residual singleton criterion). *The profile F reconstructs τ if and only if $R_F^\tau(x)$ is a singleton for every $x \in \mathcal{C}$.*

Proof. If F reconstructs τ , then all points in $E_F[x]$ have the same τ -value as x , so $R_F^\tau(x) = \{\tau(x)\}$. Conversely, if every residual target set is a singleton and $x E_F y$, then $y \in E_F[x]$, so $\tau(y) = \tau(x)$. Thus $E_F \subseteq K_\tau$. $\square$

Definition 7.6 (Bad pairs). The bad-pair set for F relative to τ is

$$\text{Bad}_F^\tau = \{(x, y) : x E_F y, \tau(x) \neq \tau(y)\}.$$

Then F reconstructs τ if and only if $\text{Bad}_F^\tau = \emptyset$.

Definition 7.7 (Limiting indistinguishability). For a filtered observable system, define the limiting indistinguishability relation

$$E_\infty = \bigcap_{b \in B} E_b.$$

Proposition 7.8 (Irremovable indistinguishability). *If $E_\infty \not\subseteq K_\tau$, then no finite budget in the declared filtration reconstructs τ . In particular, every window reconstruction threshold for τ is infinite.*

Proof. Choose x, y with $x E_\infty y$ and $\tau(x) \neq \tau(y)$. Since $E_\infty \subseteq E_b$ for every b , the pair x, y is indistinguishable at every finite budget. Hence no E_b is contained in K_τ . $\square$

7.1 Residual ambiguity profiles

The residual set $R_F^\tau(x)$ records which target values remain possible after the F -profile of x has been read. Different applications need different coarse measurements of this residual ambiguity.

Definition 7.9 (Combinatorial ambiguity). Assume the residual target sets are finite. The maximum target ambiguity left by F is

$$A_F^\tau = \sup_{x \in \mathcal{C}} |R_F^\tau(x)|.$$

Thus $A_F^\tau = 1$ means that every observational fibre contains at most one target value.

Definition 7.10 (Residual entropy). Suppose $\mathcal{C}$ is equipped with a probability measure or probability distribution μ , and suppose the random variables τ and Φ_F are measurable. The residual entropy of the target after observing F is

$$H_\mu(\tau \mid \Phi_F).$$

When the target is discrete, this is the usual conditional entropy of the target value given the observational profile.

Definition 7.11 (Residual diameter). Suppose the target space T has a metric d_T . The residual target diameter left by F is

$$D_F^\tau = \sup_{x \in \mathcal{C}} \text{diam}_T R_F^\tau(x).$$

More generally, for a budget b , write $D_b^\tau = D_{\Omega_{\leq b}}^\tau$.

Proposition 7.12 (Residual criteria). *The following implications hold.*

- (a) *If all residual target sets are finite, then F reconstructs τ if and only if $A_F^\tau = 1$.*
- (b) *If T is a metric space in which distinct target values have positive distance from each other, then F reconstructs τ if $D_F^\tau = 0$.*
- (c) *If τ is discrete and $H_\mu(\tau \mid \Phi_F) = 0$, then τ is determined by Φ_F for μ -almost every object.*

Proof. For (a), $A_F^\tau = 1$ means that every set $R_F^\tau(x)$ is a singleton, and this is exactly the residual singleton criterion.

For (b), $D_F^\tau = 0$ means that each residual target set has metric diameter zero. Under the stated separation hypothesis on T , a subset of diameter zero contains at most one target value. Hence every residual target set is a singleton, so F reconstructs τ .

For (c), the standard property of conditional entropy for a discrete random variable says that $H_\mu(\tau \mid \Phi_F) = 0$ precisely when τ is almost surely a measurable function of Φ_F . Equivalently, outside a μ -null set, each observed profile determines a single target value. $\square$

8 Relative reconstruction and resolving refinements

Definition 8.1 (Relative reconstruction). Let $\sigma : \mathcal{C} \rightarrow S$ and $\tau : \mathcal{C} \rightarrow T$ be targets. A profile $F \subseteq \Omega$ reconstructs τ relative to σ if

$$E_F \cap K_\sigma \subseteq K_\tau.$$

Equivalently,

$$\sigma(x) = \sigma(y) \text{ and } \Phi_F(x) = \Phi_F(y) \implies \tau(x) = \tau(y).$$

Relative reconstruction is the natural language when some base information has already been fixed: a rank invariant, a local profile, a principal Hecke profile, a Weisfeiler–Leman profile, a semisimplification, or a coarse quotient.

Definition 8.2 (Resolving refinement). Let $\Sigma \subseteq \Omega$ be a base profile and let τ be a target. A family $\Lambda \subseteq \Omega$ is a resolving refinement for τ over Σ if

$$E_\Sigma \cap E_\Lambda \subseteq K_\tau.$$

Remark 8.3. Additional observables refine the base quotient until all target-distinct pairs in the residual fibres have been separated.

9 Cost models

9.1 Window thresholds

Assume the observable system is filtered by a totally ordered budget scale. Define the window reconstruction threshold by

$$W_\Omega(\tau) = \inf\{b : E_b \subseteq K_\tau\}.$$

Similarly, the relative window threshold is

$$W_\Omega(\tau \mid \sigma) = \inf\{b : E_b \cap K_\sigma \subseteq K_\tau\}.$$

Definition 9.1 (Pair separation cost). For $x, y \in \mathcal{C}$, define

$$s_\Omega(x, y) = \inf\{\chi(\omega) : \omega(x) \neq \omega(y)\}.$$

If no allowed observable separates x and y , set $s_\Omega(x, y) = \infty$.

Proposition 9.2 (Pair formula for window thresholds). *Assume the budget scale is ordered so that $\Omega_{\leq b}$ has the expected meaning. Then*

$$W_\Omega(\tau) = \sup_{\tau(x) \neq \tau(y)} s_\Omega(x, y),$$

with the usual convention that the equality holds in the extended ordered scale. More generally,

$$W_\Omega(\tau \mid \sigma) = \sup_{\substack{\sigma(x) = \sigma(y) \\ \tau(x) \neq \tau(y)}} s_\Omega(x, y).$$

Proof. A budget b reconstructs τ exactly when

$$E_b \subseteq K_\tau.$$

Unwinding the definitions, this says that there is no pair x, y with $\tau(x) \neq \tau(y)$ and $x E_b y$. Equivalently, for every pair with different target values, at least one observable in $\Omega_{\leq b}$ separates the pair.

For a fixed target-distinct pair x, y , the condition that some observable of cost at most b separates the pair is precisely

$$s_\Omega(x, y) \leq b.$$

Thus b reconstructs τ if and only if b is an upper bound for all numbers $s_\Omega(x, y)$ with $\tau(x) \neq \tau(y)$. The least such budget is therefore their supremum.

The relative statement is the same argument applied only to pairs that remain in the same base target fibre, namely pairs satisfying $\sigma(x) = \sigma(y)$. Such pairs are exactly the pairs on which K_σ imposes the residual reconstruction problem. $\square$

9.2 Multi-resource regions

Many observation systems have several simultaneous costs. The relevant object is then a region of feasible budgets.

Definition 9.3 (Reconstruction region). Assume B is a partially ordered budget set and the filtration is monotone: $b \leq b'$ implies $\Omega_{\leq b} \subseteq \Omega_{\leq b'}$. Define

$$\mathcal{R}_\Omega(\tau) = \{b \in B : E_b \subseteq K_\tau\}.$$

The relative reconstruction region is

$$\mathcal{R}_\Omega(\tau \mid \sigma) = \{b \in B : E_b \cap K_\sigma \subseteq K_\tau\}.$$

Proposition 9.4 (Upward closure). *The regions $\mathcal{R}_\Omega(\tau)$ and $\mathcal{R}_\Omega(\tau \mid \sigma)$ are upward closed in B .*

Proof. Suppose $b \in \mathcal{R}_\Omega(\tau)$ and $b \leq b'$. Monotonicity gives $\Omega_{\leq b} \subseteq \Omega_{\leq b'}$, hence

$$E_{b'} \subseteq E_b.$$

Since $E_b \subseteq K_\tau$, it follows that $E_{b'} \subseteq K_\tau$, so $b' \in \mathcal{R}_\Omega(\tau)$.

For the relative region, the same inclusion gives

$$E_{b'} \cap K_\sigma \subseteq E_b \cap K_\sigma \subseteq K_\tau.$$

Thus $b' \in \mathcal{R}_\Omega(\tau \mid \sigma)$. $\square$

Definition 9.5 (Pareto frontier). The Pareto frontier of a reconstruction region is the set of minimal elements of that region, when they exist. It records budgets that reconstruct the target and cannot be improved in one resource without worsening another or leaving the region.

9.3 Finite fingerprints

In the finite-fingerprint model, one chooses a finite family of observables and pays an aggregate cost. For additive costs define

$$M_\Omega(\tau) = \inf \left\{ \sum_{\omega \in F} \chi(\omega) : F \subseteq \Omega \text{ finite and } E_F \subseteq K_\tau \right\}.$$

The relative version over a base profile Σ is

$$M_\Omega(\tau \mid \Sigma) = \inf \left\{ \sum_{\omega \in F} \chi(\omega) : F \subseteq \Omega \text{ finite and } E_\Sigma \cap E_F \subseteq K_\tau \right\}.$$

9.4 Adaptive reconstruction

An adaptive strategy chooses observables sequentially, with later choices depending on earlier values. Formally, such a strategy is a decision tree whose internal nodes are observables and whose outgoing edges are possible observed values.

Definition 9.6 (Adaptive reconstruction). An adaptive strategy reconstructs τ if every leaf of the decision tree is target-homogeneous: all objects compatible with the transcript at that leaf have the same τ -value. Its worst-case cost is the maximum aggregate cost along a root-to-leaf path.

Proposition 9.7 (Adaptive reconstruction equals decision-tree reconstruction). *An adaptive strategy reconstructs τ if and only if every leaf of its decision tree is contained in a fibre of τ .*

Proof. The transcript at a leaf records exactly the observable values encountered along the corresponding path. The objects reaching that leaf are precisely the objects compatible with those values. The transcript determines the target exactly when this compatible set has a single target value. This is the leaf-homogeneity condition. $\square$

10 Finite reconstruction as set cover

Assume now that $\mathcal{C}$ is finite and that the target $\tau : \mathcal{C} \rightarrow T$ is fixed.

Definition 10.1 (Target-distinct pairs). Let

$$P_\tau = \{\{x, y\} : x, y \in \mathcal{C}, \tau(x) \neq \tau(y)\}$$

be the set of unordered target-distinct pairs. For an observable $\omega \in \Omega$, define

$$S_\omega = \{\{x, y\} \in P_\tau : \omega(x) \neq \omega(y)\}.$$

Theorem 10.2 (Finite fingerprint equals weighted set cover). *Let $\mathcal{C}$ be finite. A finite family $F \subseteq \Omega$ reconstructs τ if and only if*

$$P_\tau = \bigcup_{\omega \in F} S_\omega.$$

Consequently, the minimum-cost finite fingerprint problem is exactly weighted set cover on the universe P_τ , with sets S_ω and weights $\chi(\omega)$.

Proof. If F reconstructs τ , let $\{x, y\} \in P_\tau$. Since $\tau(x) \neq \tau(y)$, the two objects cannot have the same F -profile. Thus some $\omega \in F$ satisfies $\omega(x) \neq \omega(y)$, and $\{x, y\} \in S_\omega$. Hence the sets S_ω cover P_τ .

Conversely, if the sets S_ω cover P_τ , then every target-distinct pair is separated by at least one observable in F . Therefore no pair in E_F has different target values, so $E_F \subseteq K_\tau$.

The optimization statement is the same equivalence with weights included. $\square$

Definition 10.3 (Residual pairs over a base profile). For a base profile Σ , define

$$P_{\tau|\Sigma} = \{\{x, y\} : x E_\Sigma y, \tau(x) \neq \tau(y)\}.$$

For a candidate observable λ , define

$$S_\lambda^\Sigma = \{\{x, y\} \in P_{\tau|\Sigma} : \lambda(x) \neq \lambda(y)\}.$$

Corollary 10.4 (Finite resolving refinements). *A finite family Λ resolves τ over Σ if and only if*

$$P_{\tau|\Sigma} = \bigcup_{\lambda \in \Lambda} S_{\lambda}^{\Sigma}.$$

Thus minimum-cost finite resolving refinement is weighted set cover on residual target-distinct pairs.

Proof. Assume first that Λ resolves τ over Σ . Let $\{x, y\} \in P_{\tau|\Sigma}$. Then $x E_{\Sigma} y$ and $\tau(x) \neq \tau(y)$. Since Λ resolves the residual ambiguity, one cannot also have $x E_{\Lambda} y$. Therefore some $\lambda \in \Lambda$ satisfies $\lambda(x) \neq \lambda(y)$, so $\{x, y\} \in S_{\lambda}^{\Sigma}$. Hence the sets S_{λ}^{Σ} cover $P_{\tau|\Sigma}$.

Conversely, suppose the displayed union covers $P_{\tau|\Sigma}$. If $x E_{\Sigma} y$ and $x E_{\Lambda} y$, then no $\lambda \in \Lambda$ separates x from y . If $\tau(x) \neq \tau(y)$, then $\{x, y\} \in P_{\tau|\Sigma}$, so the covering assumption would put $\{x, y\}$ in some S_{λ}^{Σ} , a contradiction. Thus $\tau(x) = \tau(y)$, proving

$$E_{\Sigma} \cap E_{\Lambda} \subseteq K_{\tau}.$$

The cost-minimization statement is again exactly the same covering problem with weights. $\square$

10.1 Dual lower bounds

The set-cover formulation gives a general lower-bound mechanism. The linear relaxation is

$$\min \sum_{\omega \in \Omega} \chi(\omega) z_{\omega}$$

subject to

$$\sum_{\omega: p \in S_{\omega}} z_{\omega} \geq 1 \quad \text{for all } p \in P_{\tau}, \quad z_{\omega} \geq 0.$$

The dual is

$$\max \sum_{p \in P_{\tau}} u_p$$

subject to

$$\sum_{p \in S_{\omega}} u_p \leq \chi(\omega) \quad \text{for all } \omega \in \Omega, \quad u_p \geq 0.$$

Principle 10.5 (Cover-packing duality). An upper bound is a family of observables covering all target-distinct pairs. A lower bound is a fractional packing of target-distinct pairs such that no cheap observable covers too much packed weight.

Theorem 10.6 (Fractional packing lower bound). *Let $u_p \geq 0$ be weights on target-distinct pairs $p \in P_{\tau}$. Suppose that for every observable $\omega \in \Omega$,*

$$\sum_{p \in S_{\omega}} u_p \leq \chi(\omega).$$

Then every finite fingerprint F reconstructing τ satisfies

$$\sum_{\omega \in F} \chi(\omega) \geq \sum_{p \in P_{\tau}} u_p.$$

Proof. Let F reconstruct τ . By the set-cover theorem, the sets S_{ω} , for $\omega \in F$, cover P_{τ} . Therefore

$$\sum_{p \in P_{\tau}} u_p \leq \sum_{\omega \in F} \sum_{p \in S_{\omega}} u_p.$$

The inequality may be strict because a pair can be covered by several observables, so it can appear more than once on the right. By the assumed packing constraint,

$$\sum_{p \in S_\omega} u_p \leq \chi(\omega)$$

for each ω . Summing over $\omega \in F$ gives

$$\sum_{p \in P_\tau} u_p \leq \sum_{\omega \in F} \chi(\omega).$$

This proves the lower bound. $\square$

This is often the cleanest way to state finite optimality or quasi-optimality for a selected family of tests. The relative version is obtained by replacing P_τ by $P_{\tau|\Sigma}$.

11 Structured residual regimes

Finite set cover is the unstructured model. In many mathematical situations the residual ambiguity has algebraic structure. Then resolving refinements are governed by familiar algebraic criteria.

11.1 Linear residuals

Proposition 11.1 (Linear residual criterion). *Let V be a finite-dimensional vector space over a field k . Suppose that, inside a fixed base fibre, target ambiguity is parametrized by V , and candidate resolving observables are linear functionals $\lambda \in L \subseteq V^*$. A family $\Lambda \subseteq L$ resolves the residual target ambiguity if and only if*

$$\bigcap_{\lambda \in \Lambda} \ker \lambda = \{0\}.$$

Equivalently, the span of Λ is all of V^ on the residual subspace. If all functionals in V^* are available at unit cost, the minimum number of probes is $\dim V$.*

Proof. Two residual parameters $v, w \in V$ are indistinguishable by Λ exactly when $\lambda(v - w) = 0$ for every $\lambda \in \Lambda$, that is, when $v - w \in \bigcap_{\lambda \in \Lambda} \ker \lambda$. The family separates all distinct residual parameters precisely when this intersection is $\{0\}$. This is equivalent to the linear map

$$V \rightarrow k^\Lambda, \quad v \mapsto (\lambda(v))_{\lambda \in \Lambda}$$

being injective, hence to Λ spanning V^* on the residual space. If all functionals are available at unit cost, at least $\dim V$ are needed and a basis of V^* suffices. $\square$

11.2 Torsorial residuals

Proposition 11.2 (Character separation of finite torsors). *Let a finite abelian group A act simply transitively on a residual target fibre, and suppose candidate resolving observables are characters $\chi : A \rightarrow k^\times$. A family $\Lambda \subseteq \text{Hom}(A, k^\times)$ separates the residual fibre if and only if*

$$\bigcap_{\chi \in \Lambda} \ker \chi = \{e\}.$$

Proof. Choosing a base point identifies the residual fibre with A . Two points $a, b \in A$ have the same values under all characters in Λ exactly when $\chi(ab^{-1}) = 1$ for every $\chi \in \Lambda$, that is, when ab^{-1} lies in the displayed intersection. The family separates all points exactly when the intersection is trivial. $\square$

11.3 Algebraic residuals

Proposition 11.3 (Rational reconstruction criterion). *Let X be an irreducible variety over a field k . Let observables $f_i \in k(X)$ define a rational map*

$$\Phi : X \dashrightarrow \mathbb{A}^m,$$

and let a rational target be represented by $g \in k(X)$. On the domain where these maps are defined, g is generically determined by Φ if

$$g \in k(f_1, \dots, f_m) \subseteq k(X).$$

If $g \notin k(f_1, \dots, f_m)$, then the profile does not determine g at the level of function fields.

Proof. The rational functions determined by the profile are exactly the elements of the subfield $k(f_1, \dots, f_m)$. If g lies in this subfield, then $g = R(f_1, \dots, f_m)$ for some rational expression R , so g is determined wherever the expression is defined. If g does not lie in the subfield, then the rational map given by the f_i does not factor the rational map g at the function-field level. $\square$

11.4 Logical residuals

Let $\mathcal{C}$ be a class of finite structures and let Ω_r be the set of formulas in a declared logical fragment of rank or width at most r . A formula $\varphi \in \Omega_r$ defines an observable

$$M \mapsto 1_{M \models \varphi}.$$

Proposition 11.4 (Logical reconstruction criterion). *Let $F = \Omega_r$, and let $P \subseteq \mathcal{C}$ be a Boolean target. Then F reconstructs P exactly when P is a union of E_F -classes. Equivalently, no two structures that satisfy the same formulas in the declared fragment have different truth values for P . If $\mathcal{C}$ is finite and the fragment is closed under finite Boolean combinations on $\mathcal{C}$, this is equivalent to definability of P on $\mathcal{C}$ by a formula built from the fragment.*

Proof. The Boolean target P has target kernel

$$K_P = (P \times P) \cup ((\mathcal{C} \setminus P) \times (\mathcal{C} \setminus P)).$$

Thus $E_F \subseteq K_P$ precisely says that every E_F -class lies entirely inside P or entirely outside P . This is the same as saying that P is a union of E_F -classes.

If $\mathcal{C}$ is finite and the fragment is closed under finite Boolean combinations, each E_F -class can be described on $\mathcal{C}$ by requiring the finite list of truth values of formulas that separates that class from the other classes. A finite union of such class descriptions defines P . Conversely, if P is defined by a Boolean combination of observables in the fragment, then two structures with the same F -profile must have the same truth value for P . $\square$

11.5 Probabilistic observables

Some observations are noisy or sampled. A probabilistic observable is a Markov kernel

$$K_\omega(a \mid x)$$

from $\mathcal{C}$ to a measurable output space A_ω . It separates x and y in the zero-error sense if the probability measures $K_\omega(\cdot \mid x)$ and $K_\omega(\cdot \mid y)$ are different.

Definition 11.5 (Stochastic post-processing). Let $K : \mathcal{C} \rightsquigarrow A$ and $L : \mathcal{C} \rightsquigarrow B$ be probabilistic observables. We say that K Blackwell-dominates L if there is a Markov kernel M from A to B such that

$$L(\cdot \mid x) = M \circ K(\cdot \mid x)$$

for every $x \in \mathcal{C}$. In words, L can be obtained from K by stochastic post-processing.

Proposition 11.6 (Post-processing cannot improve zero-error separation). *If K Blackwell-dominates L , then every pair separated by L is separated by K .*

Proof. Suppose $K(\cdot | x) = K(\cdot | y)$. Applying the same post-processing kernel M to both measures gives

$$L(\cdot | x) = M \circ K(\cdot | x) = M \circ K(\cdot | y) = L(\cdot | y).$$

Thus L cannot separate a pair that K fails to separate. Equivalently, stochastic post-processing can only coarsen the zero-error indistinguishability relation. $\square$

11.6 Local-global residuals

Local observations often reconstruct local targets while leaving gluing ambiguity. The following standard descent pattern is useful when the hypotheses can be verified in a concrete category.

Proposition 11.7 (Local-global residual pattern). *Let an object be reconstructed locally on a cover $\{U_i\}$. Suppose:*

- (a) *the chosen local observables determine the local target objects on every U_i ;*
- (b) *every remaining ambiguity in gluing the local reconstructions is represented by a sheaf or presheaf $\mathcal{A}$ of invisible local automorphisms;*
- (c) *compatible global reconstructions with the same local profiles are classified by the corresponding descent cocycles.*

Then the residual global ambiguity is governed by the nonabelian cohomology set $H^1(\{U_i\}, \mathcal{A})$. In the abelian case this is an ordinary cohomology group. If the relevant H^1 is trivial, the local reconstruction has no residual gluing ambiguity.

Proof. Choose local reconstructed objects and local identifications on overlaps. Changing the local choices by invisible automorphisms changes the overlap identifications by a 1-coboundary. The compatibility condition on triple overlaps is the cocycle condition. Therefore equivalence classes of gluings with the same local observed data are represented by 1-cocycles modulo 1-coboundaries, which is precisely the stated H^1 classification in the standard descent formalism. If this cohomology set is trivial, every compatible local reconstruction is equivalent to the same global reconstruction. $\square$

12 Stability

Exact reconstruction is set-theoretic. In analytic, geometric, statistical, and numerical settings one also needs stability.

Definition 12.1 (Reconstruction modulus). Assume the profile space has a metric d_P and the target space has a metric d_T . Define

$$\alpha_{F,\tau}(\varepsilon) = \sup\{d_T(\tau(x), \tau(y)) : d_P(\Phi_F(x), \Phi_F(y)) \leq \varepsilon\}.$$

Definition 12.2 (Stable reconstruction). The profile F stably reconstructs τ if

$$\alpha_{F,\tau}(\varepsilon) \rightarrow 0 \quad \text{as} \quad \varepsilon \rightarrow 0.$$

Exact reconstruction is the condition $\alpha_{F,\tau}(0) = 0$.

Proposition 12.3 (Approximate twins obstruct stability). *Suppose there are sequences $x_n, y_n \in \mathcal{C}$ and a number $\delta > 0$ such that*

$$d_P(\Phi_F(x_n), \Phi_F(y_n)) \rightarrow 0$$

but

$$d_T(\tau(x_n), \tau(y_n)) \geq \delta$$

for all n . Then F does not stably reconstruct τ .

Proof. For every sufficiently small $\varepsilon > 0$, some pair x_n, y_n satisfies $d_P(\Phi_F(x_n), \Phi_F(y_n)) \leq \varepsilon$ while the target distance is at least δ . Therefore $\alpha_{F, \tau}(\varepsilon) \geq \delta$ along a sequence $\varepsilon \rightarrow 0$, so the modulus does not tend to zero. $\square$

Proposition 12.4 (Infinitesimal obstruction). *Let X, P, T be smooth manifolds, and let*

$$\Phi : X \rightarrow P, \quad \tau : X \rightarrow T$$

be smooth maps. If τ locally factors through Φ near x , then

$$\ker d\Phi_x \subseteq \ker d\tau_x.$$

Proof. If $\tau = R \circ \Phi$ locally, then

$$d\tau_x = dR_{\Phi(x)} \circ d\Phi_x.$$

Every tangent vector killed by $d\Phi_x$ is therefore killed by $d\tau_x$. $\square$

13 Transfer and verifiable completeness

Observable reconstruction transfers only when observable access is controlled. The following simple form is often enough.

Theorem 13.1 (Transfer of window upper bounds). *Let $F : \mathcal{C} \rightarrow \mathcal{D}$ be a map of object classes. Suppose observable systems on $\mathcal{C}$ and $\mathcal{D}$ have budget relations $E_b^{\mathcal{C}}$ and $E_b^{\mathcal{D}}$. Assume there is an overhead function a such that*

$$xE_{a(b)}^{\mathcal{C}}y \implies F(x)E_b^{\mathcal{D}}F(y)$$

for all $x, y \in \mathcal{C}$. If $\Omega_{\mathcal{D}}$ reconstructs a target $\tau_{\mathcal{D}}$ at budget b , then $\Omega_{\mathcal{C}}$ reconstructs $\tau_{\mathcal{D}} \circ F$ at budget $a(b)$.

Proof. Assume $xE_{a(b)}^{\mathcal{C}}y$. By the overhead hypothesis, $F(x)E_b^{\mathcal{D}}F(y)$. Since budget b reconstructs $\tau_{\mathcal{D}}$, one has

$$\tau_{\mathcal{D}}(F(x)) = \tau_{\mathcal{D}}(F(y)).$$

Thus $E_{a(b)}^{\mathcal{C}} \subseteq K_{\tau_{\mathcal{D}} \circ F}$, which is the desired reconstruction statement. $\square$

For verifiable access, a proof of observational completeness is a proof of an inclusion

$$E_F \subseteq K_{\tau}$$

or, in the relative case,

$$E_{\Sigma} \cap E_F \subseteq K_{\tau}.$$

Definition 13.2 (Completeness proof data). Fix a declared verification model. Completeness proof data for a profile F relative to a target τ is finite auxiliary data which a checker can use to verify the inclusion

$$E_F \subseteq K_\tau.$$

Relative completeness proof data for F over a base profile Σ verifies

$$E_\Sigma \cap E_F \subseteq K_\tau.$$

Proposition 13.3 (Standard forms of completeness proof data). *The following are standard forms of completeness proof data, once the relevant computations are made auditable in the declared verification model.*

- (a) *In a finite object class, a cover of all target-distinct pairs by separating observables proves $E_F \subseteq K_\tau$.*
- (b) *In a linear residual problem, a full-rank matrix of selected linear functionals proves that the residual kernel is zero.*
- (c) *In a torsorial residual problem, a computation of $\bigcap_i \ker \chi_i = \{e\}$ proves that selected characters separate the residual torsor.*
- (d) *In an algebraic residual problem, a field-generation identity or an ideal-membership argument can prove that the target function is determined by the observed functions on the declared locus.*

Proof. In (a), the finite set-cover theorem says that covering all target-distinct pairs is equivalent to the inclusion $E_F \subseteq K_\tau$. In (b), the linear residual criterion says that zero common kernel is equivalent to separating all residual parameters. In (c), the torsorial criterion gives the same equivalence for characters on the residual group. In (d), if the target function lies in the field generated by the observed functions, or if the corresponding graph inclusion is proved by ideal membership on the declared locus, then equal observed values force equal target values. Each item is therefore a concrete proof of the relevant equivalence-relation inclusion. $\square$

This connects the observable calculus here with the audit layer of Presentation Theory: the mathematical assertion is the inclusion of equivalence relations, while the audit problem asks which bounded checker can verify the supplied proof data.

14 Three model applications

The applications below are schematic: their purpose is to show how the same quotient calculus appears in different domains.

14.1 Rank invariants in multiparameter persistence

Let $\mathcal{C}$ be a class of multiparameter persistence modules, and let F be the family of rank probes

$$\omega_{a,b}(M) = \text{rank}(M(a) \rightarrow M(b)).$$

The rank profile induces E_F . If the target is isomorphism class, then incompleteness is the statement

$$E_F \not\subseteq K_{\text{iso}}.$$

The stronger presentation-theoretic question asks for the residual target sets

$$R_F^{\text{iso}}(M)$$

and for their geometry. In known multiparameter constructions, rank-invariant fibres can contain positive-dimensional projective moduli and, in richer Boolean-lattice versions, realization spaces with arbitrary algebraic singularity type. In the language of this paper, the rank profile determines the quotient $\mathcal{C}/E_F$, while the classification problem remains inside residual fibres of that quotient.

14.2 Principal Hecke data and class-group twists

Let $\mathcal{C}$ be a finite computational window of Hecke eigensystems over a number field K . Let Σ be the profile of principal Hecke data. If unramified quadratic twists by characters of $\text{Cl}(K)$ preserve the principal profile, then Σ leaves a residual torsorial ambiguity.

Suppose the residual twist group is

$$X \subseteq \text{Hom}(\text{Cl}(K), \{\pm 1\}).$$

Non-principal probes evaluate these characters on ideal classes $c_1, \dots, c_m$. By the torsorial criterion, the probes resolve the residual ambiguity exactly when

$$\bigcap_{j=1}^m \ker(\text{ev}_{c_j} |_X) = \{1\}.$$

If $X \cong (\mathbb{F}_2)^d$ and all needed linear functionals are available at unit cost, d independent probes are necessary and sufficient. With weighted probes, the optimization is a minimum-cost basis problem in the corresponding linear matroid.

14.3 Finite hom-count completions

Let $\mathcal{C}$ be a finite graph class, let $\mathcal{F}_0$ be a base family of test graphs, and let the target be graph isomorphism on $\mathcal{C}$. The base profile leaves residual pairs

$$P_{\text{iso}|\mathcal{F}_0} = \{\{G, H\} : G \not\cong H, \Phi_{\mathcal{F}_0}(G) = \Phi_{\mathcal{F}_0}(H)\}.$$

For a candidate test graph T , define

$$S_T = \{\{G, H\} \in P_{\text{iso}|\mathcal{F}_0} : \#\text{Hom}(T, G) \neq \#\text{Hom}(T, H)\}.$$

A finite family $\mathcal{T}$ completes the profile exactly when

$$P_{\text{iso}|\mathcal{F}_0} = \bigcup_{T \in \mathcal{T}} S_T.$$

Thus exact finite completion by additional hom-count tests is precisely a set-cover problem on residual non-isomorphic pairs. The linear-programming dual gives a lower-bound method: assign weights to residual pairs so that every test graph of a given cost covers only bounded total weight.

15 Summary

Observable reconstruction can be summarized by four layers.

First, every profile F induces an equivalence relation E_F , hence an observable quotient $\mathcal{C}/E_F$. Every target τ induces K_τ . Reconstruction is the inclusion $E_F \subseteq K_\tau$.

Second, the separation preorder, observational closure, and the observable-equivalence Galois correspondence organize the information content of profiles. Post-processing changes representation, not separation power.

Third, incomplete profiles still determine a precise target factor:

$$Q_F(\tau) = \mathcal{C}/(K_\tau \vee E_F).$$

Relative reconstruction and resolving refinements describe how to eliminate residual target ambiguity inside a base quotient.

Fourth, costs turn the quotient calculus into optimization. Window costs are governed by the hardest target-distinct pair to separate. Finite fingerprints are weighted set cover. Adaptive reconstruction is decision-tree reconstruction. Structured residuals reduce to standard algebraic criteria. Stability is measured by a reconstruction modulus.

This is the observable layer of Presentation Theory: a common language for the quotient and cost structure shared by reconstruction arguments in invariant theory, representation theory, graph theory, persistence theory, arithmetic, and related domains.

References

- [1] D. Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics* 24 (1953), 265–272.
- [2] G. Carlsson and A. Zomorodian. The theory of multidimensional persistence. *Discrete & Computational Geometry* 42 (2009), 71–93.
- [3] J. Cai, M. Fuerer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica* 12 (1992), 389–410.
- [4] M. Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*. Cambridge University Press, 2017.
- [5] L. Lovasz. *Large Networks and Graph Limits*. American Mathematical Society, 2012.
- [6] S. Mac Lane. *Categories for the Working Mathematician*. Second edition, Springer, 1998.