

Presentation Theory II: Controlled Transfer, Observable Budgets, and the Geometry of Fibres

Luca Blanchi

June 2026

Abstract

The first note on Presentation Theory introduced presentation systems, bounded parts, realization fibres, observables, observable costs, and normal-form compilers. This sequel develops the transfer calculus. A map between mathematical worlds is useful for Presentation Theory only when it controls access: descriptions, observables, operations, verification data, and fibres. We package this structure as a transfer package.

The formal part proves reusable transfer principles. A presentation morphism sends bounded source objects to bounded target objects with controlled overhead. Its optimal distortion profile composes along paths. In effective filtered systems, equivalence reductions transfer equivalence degrees, while bounded-presentation reductions transfer bounded-presentation degrees; mere cost control gives only an upper bound and does not by itself reflect low-cost representatives. Observable-compatible morphisms pull back bounded observable images and distinguishing-cost lower bounds. Operation-compatible morphisms transport term constructions. Base-fibre decompositions split complexity into a base cost and a relative fibre cost when assembly and extraction are controlled. Verification relations transfer proof data in the directions supplied by the package. These ingredients assemble into a controlled transfer schema for universal statements in a budgeted fragment.

The sequel then isolates several engines that recur across applications. Fragmented positive predicates have optimal resource profiles of the same Turing degree as the underlying decision problem; their total majorants form exactly the upper Turing cone above that problem, giving a no-effective-atlas principle for access tasks with undecidable positive part. Finite-window systems turn undecidable global existence into non-computable obstruction radii. Undecidable threshold problems for computable limiting quantities forbid computable convergence moduli. Move-connected presentation fibres have bridge-height profiles of the same Turing degree as recognition of the fibre. Rank-profile compilers transfer Murphy-type universality into fibres of observables. Wild-fibre compilers place matrix-pair classification problems inside fixed observable fibres. A budgeted Morita principle records when two presentation contexts have the same quantitative theory up to a chosen overhead class.

The applications cover finite models, Diophantine equations, proof length, Dehn area, cellular automata, grammar ambiguity, matrix mortality, Wang tilings, nonlocal games, multiparameter persistence, rank-derived invariants, finite poset sheaves, quiver representations, finite group quotients, finite linear representations, Pachner moves, and Tietze moves. The common point is that transfer turns qualitative universality and undecidability results into quantitative statements about observable budgets and fibre geometry.

1 Purpose

Presentation Theory studies mathematical objects through access systems:

descriptions $\longrightarrow$ objects.

A presentation system tells us which descriptions are allowed, how they realize objects, and what they cost. Observables then extract information from objects, often with their own cost. Realization maps and observable maps have fibres, and those fibres measure the information forgotten by the chosen access system.

The first note developed this basic language. This note studies the next question:

when can a block of mathematics be transported from one presentation context to another?

A plain map of object classes is rarely enough. If one wants to transport lower bounds, reconstruction statements, non-computability results, or fibre geometry, the map must control descriptions, observables, operations, verification data, and fibres. The resulting object is a transfer package.

The guiding principle is:

a map transports objects;
a transfer package transports access to objects.

2 Presentation Systems

We recall the basic objects, using notation compatible with the first note.

Definition 2.1 (Resource scale). *A resource scale is a preordered set $(B, \leq)$. Examples include*

$$\mathbb{N}, \quad \mathbb{N}^k, \quad \mathbb{R}_{\geq 0},$$

with the usual or coordinatewise orders. The scale records costs such as length, degree, height, arity, number of generators, number of relations, circuit size, proof length, or a vector of several resources.

Definition 2.2 (Overhead class). *An overhead class $\mathcal{H}$ between resource scales is a collection of monotone maps between resource scales, containing identities and closed under composition. Typical choices are linear, polynomial, primitive recursive, computable, or arbitrary monotone overheads.*

Remark 2.3 (Admissible overheads). *An overhead function is part of the declared comparison data. It is not required to be optimal, but it must be compatible with the resource scales under consideration and monotone with respect to their orders. In effective settings one usually restricts to computable overheads; in quantitative complexity settings one may restrict further to linear, polynomial, or primitive recursive overheads. The distortion profile records the possible overheads for a given comparison, while a chosen function f records a specific upper bound.*

Definition 2.4 (Presentation system). *A presentation system on an object class $\mathcal{X}$ is a triple*

$$\Gamma = (\mathcal{D}, \rho, \kappa),$$

where $\mathcal{D}$ is a class of descriptions,

$$\rho : \mathcal{D} \rightarrow \mathcal{X}$$

is a realization map, and

$$\kappa : \mathcal{D} \rightarrow B$$

is a cost function.

If $\mathcal{X}$ is studied up to an equivalence relation $\equiv_{\mathcal{X}}$, then $\rho(d) = x$ means $\rho(d) \equiv_{\mathcal{X}} x$.

Definition 2.5 (Bounded parts). For $b \in B$, define

$$\mathcal{D}^{\leq b} = \{d \in \mathcal{D} : \kappa(d) \leq b\}$$

and

$$\mathcal{X}_\Gamma^{\leq b} = \{x \in \mathcal{X} : \exists d \in \mathcal{D}, \rho(d) = x, \kappa(d) \leq b\}.$$

When the infimum exists, the presentation complexity of x is

$$C_\Gamma(x) = \inf\{\kappa(d) : \rho(d) = x\}.$$

The filtration $b \mapsto \mathcal{X}_\Gamma^{\leq b}$ is often more fundamental than the numerical complexity C_Γ , especially for vector-valued or partially ordered resource scales.

3 Observable Systems

Descriptions produce objects. Observables extract information from objects.

Definition 3.1 (Observable system). An observable system on $\mathcal{X}$ is a triple

$$\mathfrak{O} = (\Omega, \chi, R),$$

where Ω is a class of observables

$$\omega : \mathcal{X} \rightarrow A_\omega,$$

$\chi : \Omega \rightarrow R$ is an observable-cost function, and R is a resource scale. For $r \in R$, write

$$\Omega^{\leq r} = \{\omega \in \Omega : \chi(\omega) \leq r\}.$$

Examples of observable cost include arity of a word-count invariant, size of a finite test object, degree of a polynomial probe, rank of a linear representation, number of samples, quantifier rank, or matrix dimension.

Definition 3.2 (Bounded observable image). Given a presentation system Γ and an observable ω , define

$$\text{Val}_{\Gamma, \omega}(b) = \omega(\mathcal{X}_\Gamma^{\leq b}) \subseteq A_\omega.$$

This is the set of observable values attained by objects with descriptions of cost at most b .

Proposition 3.3 (Bounded-image obstruction). Suppose $S_b \subseteq A_\omega$ satisfies

$$\text{Val}_{\Gamma, \omega}(b) \subseteq S_b.$$

If

$$\omega(x) \notin S_b,$$

then $x \notin \mathcal{X}_\Gamma^{\leq b}$. In particular, when C_Γ is defined, $C_\Gamma(x) \not\leq b$. In a totally ordered numerical scale this may be written $C_\Gamma(x) > b$.

Proof. If $x \in \mathcal{X}_\Gamma^{\leq b}$, then by definition x is in the bounded class. Applying ω gives

$$\omega(x) \in \omega(\mathcal{X}_\Gamma^{\leq b}) = \text{Val}_{\Gamma, \omega}(b) \subseteq S_b.$$

The stated implication is the contrapositive. □

Definition 3.4 (Distinguishing cost). For $x, x' \in \mathcal{X}$, define

$$dc_\mathfrak{O}(x, x') = \inf\{\chi(\omega) : \omega(x) \neq \omega(x')\}.$$

If no allowed observable distinguishes x and x' , set $dc_\mathfrak{O}(x, x') = \infty$.

4 Presentation Contexts

A presentation system alone does not record all structures needed for transfer.

Definition 4.1 (Presentation context). *A presentation context is a tuple*

$$\mathfrak{P} = (\mathcal{X}, \Gamma, \mathfrak{D}, \Sigma, \mathfrak{V}, \mathfrak{F}, \mathfrak{M}),$$

where:

- (i) $\mathcal{X}$ is an object class;
- (ii) $\Gamma = (\mathcal{D}, \rho, \kappa)$ is a presentation system;
- (iii) $\mathfrak{D} = (\Omega, \chi, R)$ is an observable system;
- (iv) Σ is a collection of operations or constructions on $\mathcal{X}$;
- (v) $\mathfrak{V}$ is a verification layer, when proof data are part of the context;
- (vi) $\mathfrak{F}$ is a chosen structure on realization fibres;
- (vii) $\mathfrak{M}$, when present, is a family of measures or densities on bounded classes.

Only the components actually supplied by a context are used.

Definition 4.2 (Verification relation). *For a property $P \subseteq \mathcal{X}$, a verification layer consists of data*

$$(\Pi, \lambda, R_P),$$

where Π is a class of proof or verification data, $\lambda : \Pi \rightarrow S$ is a cost, and

$$R_P(x, \pi)$$

is a declared verification relation satisfying

$$R_P(x, \pi) \implies x \in P.$$

The layer is complete on a subclass $P_0 \subseteq P$ if every $x \in P_0$ admits some π with $R_P(x, \pi)$. The verification cost of $x \in P$ is

$$v_P(x) = \inf\{\lambda(\pi) : R_P(x, \pi)\}.$$

If the displayed set is empty, $v_P(x)$ is interpreted as ∞ in the extended resource scale. All overhead functions used with verification costs are understood to extend by $h(\infty) = \infty$.

Definition 4.3 (Realization fibre). *For $\Gamma = (\mathcal{D}, \rho, \kappa)$, the fibre over $x \in \mathcal{X}$ is*

$$\mathcal{F}_\rho(x) = \{d \in \mathcal{D} : \rho(d) = x\}.$$

If a move system is specified on $\mathcal{D}$, each fibre becomes a graph. Its geometry measures the cost of moving between descriptions of the same object.

5 Presentation Morphisms

Definition 5.1 (Presentation morphism). *Let*

$$\Gamma_{\mathcal{X}} = (\mathcal{D}_{\mathcal{X}}, \rho_{\mathcal{X}}, \kappa_{\mathcal{X}})$$

and

$$\Gamma_{\mathcal{Y}} = (\mathcal{D}_{\mathcal{Y}}, \rho_{\mathcal{Y}}, \kappa_{\mathcal{Y}})$$

be presentation systems. A presentation morphism with overhead f is a pair

$$(F, \hat{F}) : \Gamma_{\mathcal{X}} \rightarrow \Gamma_{\mathcal{Y}},$$

where $F : \mathcal{X} \rightarrow \mathcal{Y}$ is a map on objects and

$$\hat{F} : \mathcal{D}_{\mathcal{X}} \rightarrow \mathcal{D}_{\mathcal{Y}}$$

is a compiler satisfying

$$\rho_{\mathcal{Y}}(\hat{F}(d)) = F(\rho_{\mathcal{X}}(d))$$

and

$$\kappa_{\mathcal{Y}}(\hat{F}(d)) \leq f(\kappa_{\mathcal{X}}(d)).$$

Here f is a monotone map from the source resource scale to the target resource scale.

Theorem 5.2 (Elementary transfer of bounded parts). *If*

$$(F, \hat{F}) : \Gamma_{\mathcal{X}} \rightarrow \Gamma_{\mathcal{Y}}$$

is a presentation morphism with overhead f , then

$$F(\mathcal{X}_{\Gamma_{\mathcal{X}}}^{\leq b}) \subseteq \mathcal{Y}_{\Gamma_{\mathcal{Y}}}^{\leq f(b)}.$$

Consequently,

$$C_{\Gamma_{\mathcal{Y}}}(F(x)) \leq f(C_{\Gamma_{\mathcal{X}}}(x))$$

whenever the displayed quantities are defined and either minimum costs are attained, or f is compatible with the relevant infimum. In particular, the statement holds in the usual discrete well-ordered cost scales.

Proof. Let $x \in \mathcal{X}_{\Gamma_{\mathcal{X}}}^{\leq b}$. Choose $d \in \mathcal{D}_{\mathcal{X}}$ such that

$$\rho_{\mathcal{X}}(d) = x, \quad \kappa_{\mathcal{X}}(d) \leq b.$$

The compiler gives

$$\rho_{\mathcal{Y}}(\hat{F}(d)) = F(x)$$

and

$$\kappa_{\mathcal{Y}}(\hat{F}(d)) \leq f(\kappa_{\mathcal{X}}(d)) \leq f(b).$$

Thus $F(x) \in \mathcal{Y}_{\Gamma_{\mathcal{Y}}}^{\leq f(b)}$.

For the complexity inequality in a discrete well-ordered scale, apply the first part to a minimum-cost description of x . In the infimum formulation, the same conclusion follows under the stated compatibility of f with the relevant infimum. $\square$

Corollary 5.3 (Reflection of lower bounds). *Assume a totally ordered numerical resource scale, or read $>$ below as “not $\leq$ ” in the declared preorder. If*

$$C_{\Gamma_{\mathcal{Y}}}(F(x)) > f(b),$$

then

$$C_{\Gamma_{\mathcal{X}}}(x) > b.$$

6 Distortion Profiles

The declared overhead of a compiler need not be optimal.

Definition 6.1 (Distortion upper set). *Let $F : \mathcal{X} \rightarrow \mathcal{Y}$ be a map between object classes equipped with presentation systems Γ and Δ . Define*

$$\text{Dist}_{F,\Gamma \rightarrow \Delta}(b) = \{c : F(\mathcal{X}_\Gamma^{\leq b}) \subseteq \mathcal{Y}_\Delta^{\leq c}\}.$$

When this upper set has a least element, write it as

$$|F|_{\Gamma \rightarrow \Delta}(b).$$

Proposition 6.2 (Minimality). *If F is realized by a compiler with overhead f , then*

$$f(b) \in \text{Dist}_{F,\Gamma \rightarrow \Delta}(b)$$

for every b . If the least element exists, then

$$|F|_{\Gamma \rightarrow \Delta}(b) \leq f(b).$$

Proof. This is exactly the bounded-part transfer theorem. □

Theorem 6.3 (Composition of distortion). *Let*

$$\mathcal{X} \xrightarrow{F} \mathcal{Y} \xrightarrow{G} \mathcal{Z}$$

be maps between presentation systems Γ, Δ, Θ . If

$$c \in \text{Dist}_{F,\Gamma \rightarrow \Delta}(b)$$

and

$$d \in \text{Dist}_{G,\Delta \rightarrow \Theta}(c),$$

then

$$d \in \text{Dist}_{G \circ F, \Gamma \rightarrow \Theta}(b).$$

In particular, when least elements exist,

$$|G \circ F|_{\Gamma \rightarrow \Theta}(b) \leq |G|_{\Delta \rightarrow \Theta}(|F|_{\Gamma \rightarrow \Delta}(b)).$$

If F and G are realized by compilers with overheads f and g , then $G \circ F$ is realized by the composed compiler with overhead $g \circ f$.

Proof. Take $x \in \mathcal{X}_\Gamma^{\leq b}$. Since c controls F ,

$$F(x) \in \mathcal{Y}_\Delta^{\leq c}.$$

Since d controls G on the c -bounded target class,

$$G(F(x)) \in \mathcal{Z}_\Theta^{\leq d}.$$

This proves the upper-set statement. The least-element inequality follows by substituting the least available c and then the least available d .

For compilers, define

$$\widehat{G \circ F} = \widehat{G} \circ \widehat{F}.$$

Then

$$\rho_{\mathcal{Z}}(\widehat{G}(\widehat{F}(d))) = G(F(\rho_{\mathcal{X}}(d)))$$

and

$$\kappa_{\mathcal{Z}}(\widehat{G}(\widehat{F}(d))) \leq g(\kappa_{\mathcal{Y}}(\widehat{F}(d))) \leq g(f(\kappa_{\mathcal{X}}(d))).$$

□

7 Controlled Simulations

Definition 7.1 (Controlled simulation). *A presentation context $\mathfrak{P}_X$ is simulated by $\mathfrak{P}_Y$ with overhead f if there exists a presentation morphism*

$$\Gamma_X \rightarrow \Gamma_Y$$

with overhead f .

Definition 7.2 ($\mathcal{H}$ -equivalence). *Let $\mathcal{H}$ be an overhead class. Two presentation systems Γ and Δ are $\mathcal{H}$ -equivalent if there are presentation morphisms*

$$F : \Gamma \rightarrow \Delta, \quad G : \Delta \rightarrow \Gamma$$

with overheads in $\mathcal{H}$, such that

$$G(F(x)) \equiv x, \quad F(G(y)) \equiv y$$

on the relevant object classes, or up to declared controlled fibre data.

Theorem 7.3 (Coarse invariance). *If Γ and Δ are $\mathcal{H}$ -equivalent through maps F, G with overheads $f, g \in \mathcal{H}$, then the corresponding complexity functions compare through f and g . In particular, growth properties stable under $\mathcal{H}$ -distortion are preserved.*

Proof. The bounded-part transfer theorem gives

$$C_\Delta(F(x)) \leq f(C_\Gamma(x))$$

and

$$C_\Gamma(G(y)) \leq g(C_\Delta(y)).$$

If $G(F(x)) \equiv x$, then

$$C_\Gamma(x) = C_\Gamma(G(F(x))) \leq g(C_\Delta(F(x))).$$

Thus C_Γ is controlled by $C_\Delta \circ F$, and the reverse comparison is analogous. $\square$

8 Effective Reductions of Filtered Systems

The preceding definitions control bounded images. In computability questions one often needs a stronger, reflected comparison between the actual decision problems attached to two filtered systems.

Let

$$\Gamma = (D_\Gamma, \kappa_\Gamma, (E_t^\Gamma)) \quad \text{and} \quad \Delta = (D_\Delta, \kappa_\Delta, (E_t^\Delta))$$

be filtered effective presentation systems in the sense of Part I.

Definition 8.1 (Equivalence reduction). *A computable map*

$$F : D_\Gamma \rightarrow D_\Delta$$

is an equivalence reduction if, for all $d, e \in D_\Gamma$,

$$d E^\Gamma e \iff F(d) E^\Delta F(e).$$

Proposition 8.2 (Transfer of equivalence degree). *If there is an equivalence reduction $F : \Gamma \rightarrow \Delta$, then*

$$EQ_\Gamma \leq_m EQ_\Delta.$$

In particular,

$$\deg_T(EQ_\Gamma) \leq \deg_T(EQ_\Delta).$$

Proof. Map the input pair (d, e) to $(F(d), F(e))$. The defining equivalence of an equivalence reduction says exactly that this many-one reduction preserves and reflects membership in the equivalence problem. $\square$

Definition 8.3 (Bounded-presentation reduction). *A bounded-presentation reduction from Γ to Δ is a pair (F, φ) , where $F : D_\Gamma \rightarrow D_\Delta$ and*

$$\varphi : D_\Gamma \times \mathbb{N} \rightarrow \mathbb{N}$$

are computable, such that for every $d \in D_\Gamma$ and every $b \in \mathbb{N}$,

$$C_\Gamma(d) \leq b \iff C_\Delta(F(d)) \leq \varphi(d, b).$$

Proposition 8.4 (Transfer of bounded-presentation degree). *If (F, φ) is a bounded-presentation reduction from Γ to Δ , then*

$$P_\Gamma \leq_m P_\Delta.$$

In particular,

$$\deg_T(P_\Gamma) \leq \deg_T(P_\Delta).$$

Proof. The computable map

$$(d, b) \mapsto (F(d), \varphi(d, b))$$

preserves and reflects membership in the bounded-presentation problem by definition. $\square$

Remark 8.5 (Cost control is one-sided). *A computable map $F : D_\Gamma \rightarrow D_\Delta$ may satisfy an upper bound*

$$\kappa_\Delta(F(d)) \leq a(\kappa_\Gamma(d), |d|)$$

for a computable overhead a . This proves that images of cheap source descriptions are cheap in the target. It does not prove that every cheap representative of $F(d)$ in the target reflects to a cheap representative of d in the source. For degree monotonicity of bounded presentation, the reflected condition in the definition of bounded-presentation reduction is the relevant hypothesis.

Proposition 8.6 (Pullback of effective invariant separation). *Let $F : D_\Gamma \rightarrow D_\Delta$ be an equivalence reduction, and let $\mathcal{J} = (J_\alpha)$ be an effective invariant family for Δ . Then*

$$F^*\mathcal{J} = (J_\alpha \circ F)$$

is an effective invariant family for Γ . If $\mathcal{J}$ separates non-equivalent target pairs up to a budget s , then $F^\mathcal{J}$ separates the corresponding source pairs up to the same invariant budget.*

Proof. Effectivity is preserved by composition with the computable map F . If $d E^\Gamma e$, then $F(d) E^\Delta F(e)$, so each $J_\alpha \circ F$ is constant on E^Γ -classes. If a target invariant separates $F(d)$ and $F(e)$, then its pullback separates d and e . $\square$

9 Observable-Compatible Transfer

Presentation morphisms control descriptions. To transport lower bounds and indistinguishability results, one must also control observables.

Definition 9.1 (Observable-compatible morphism). *Let $F : \mathcal{X} \rightarrow \mathcal{Y}$ be a map between object classes equipped with observable systems*

$$\mathfrak{O}_\mathcal{X} = (\Omega_\mathcal{X}, \chi_\mathcal{X}, R_\mathcal{X}), \quad \mathfrak{O}_\mathcal{Y} = (\Omega_\mathcal{Y}, \chi_\mathcal{Y}, R_\mathcal{Y}).$$

We say that F is observable-compatible with overhead g if every target observable

$$\eta \in \Omega_{\mathcal{Y}}$$

pulls back to a source observable

$$F^*\eta = \eta \circ F \in \Omega_{\mathcal{X}}$$

and

$$\chi_{\mathcal{X}}(F^*\eta) \leq g(\chi_{\mathcal{Y}}(\eta)).$$

Theorem 9.2 (Observable image transfer). *Suppose $F : \Gamma_{\mathcal{X}} \rightarrow \Gamma_{\mathcal{Y}}$ is a presentation morphism with overhead f , and let η be a target observable. Then*

$$(\eta \circ F)(\mathcal{X}_{\Gamma_{\mathcal{X}}}^{\leq b}) \subseteq \eta(\mathcal{Y}_{\Gamma_{\mathcal{Y}}}^{\leq f(b)}).$$

Equivalently,

$$\text{Val}_{\Gamma_{\mathcal{X}}, \eta \circ F}(b) \subseteq \text{Val}_{\Gamma_{\mathcal{Y}}, \eta}(f(b)).$$

If F is observable-compatible, the pulled-back observable has cost controlled by g .

Proof. By bounded-part transfer,

$$F(\mathcal{X}_{\Gamma_{\mathcal{X}}}^{\leq b}) \subseteq \mathcal{Y}_{\Gamma_{\mathcal{Y}}}^{\leq f(b)}.$$

Applying η gives the inclusion. Observable compatibility supplies the cost bound for the pulled-back observable. $\square$

Corollary 9.3 (Pullback of bounded-image obstructions). *Suppose*

$$\eta(\mathcal{Y}_{\Gamma_{\mathcal{Y}}}^{\leq c}) \subseteq S_c.$$

Then

$$(\eta \circ F)(\mathcal{X}_{\Gamma_{\mathcal{X}}}^{\leq b}) \subseteq S_{f(b)}.$$

Hence

$$\eta(F(x)) \notin S_{f(b)} \implies x \notin \mathcal{X}_{\Gamma_{\mathcal{X}}}^{\leq b}.$$

Proposition 9.4 (Distinguishing-cost transfer). *Assume F is observable-compatible with overhead g . If a target observable η of cost r distinguishes $F(x)$ and $F(x')$, then a source observable of cost at most $g(r)$ distinguishes x and x' . In well-ordered cost scales, this gives*

$$dc_{\mathcal{X}}(x, x') \leq g(dc_{\mathcal{Y}}(F(x), F(x'))).$$

Proof. If $\eta(F(x)) \neq \eta(F(x'))$, then

$$(\eta \circ F)(x) \neq (\eta \circ F)(x').$$

Observable compatibility says that $\eta \circ F$ is allowed in the source at cost at most $g(r)$. Taking the least distinguishing cost gives the displayed inequality when minima exist. $\square$

Definition 9.5 (Observable density). *Let $\mathcal{C} \subseteq \mathcal{X}$. A map $F : \mathcal{X} \rightarrow \mathcal{Y}$ is observable-dense on $\mathcal{C}$ with overhead h if every source observable*

$$\omega \in \Omega_{\mathcal{X}}^{\leq r}$$

factors on $\mathcal{C}$ as

$$\omega = a \circ \eta \circ F,$$

where

$$\eta \in \Omega_{\mathcal{Y}}^{\leq h(r)}$$

and a is a post-processing map whose cost is either free by convention or separately accounted for.

Theorem 9.6 (Blindness transfer). *Assume F is observable-dense on $\mathcal{C}$ with overhead h . If no target observable of cost at most $h(r)$ distinguishes $F(x)$ from $F(x')$, then no source observable of cost at most r distinguishes x from x' , for $x, x' \in \mathcal{C}$.*

Proof. Suppose a source observable $\omega \in \Omega_{\mathcal{X}}^{\leq r}$ distinguishes x and x' . By observable density,

$$\omega = a \circ \eta \circ F$$

on $\mathcal{C}$, with $\eta \in \Omega_{\mathcal{Y}}^{\leq h(r)}$. Since a is a function, $\omega(x) \neq \omega(x')$ implies

$$\eta(F(x)) \neq \eta(F(x')).$$

This contradicts the assumed target indistinguishability. $\square$

10 Operations and Term Transfer

Many statements involve operations, not just individual objects.

Definition 10.1 (Operational presentation context). *Let Σ be a signature of operations. A presentation context is Σ -operational if each operation $\sigma \in \Sigma$, of arity n , is realized by a map*

$$\sigma_{\mathcal{X}} : \mathcal{X}^n \rightarrow \mathcal{X}$$

and by a description-level compiler

$$\hat{\sigma} : \mathcal{D}^n \rightarrow \mathcal{D}$$

with overhead

$$\kappa(\hat{\sigma}(d_1, \dots, d_n)) \leq \alpha_{\sigma}(\kappa(d_1), \dots, \kappa(d_n)).$$

Definition 10.2 (Operation-compatible morphism). *A presentation morphism $F : \mathcal{X} \rightarrow \mathcal{Y}$ is Σ -compatible if, for every $\sigma \in \Sigma$,*

$$F(\sigma_{\mathcal{X}}(x_1, \dots, x_n)) \equiv \sigma_{\mathcal{Y}}(F(x_1), \dots, F(x_n)),$$

possibly modulo a declared fibre equivalence.

Theorem 10.3 (Term transfer). *Let t be a term built from operations in Σ . If F is Σ -compatible and all operations have controlled overhead, then there is a recursively defined overhead α_t such that*

$$C_{\mathcal{Y}}(F(t_{\mathcal{X}}(x_1, \dots, x_n))) \leq \alpha_t(C_{\mathcal{X}}(x_1), \dots, C_{\mathcal{X}}(x_n)).$$

Proof. Induct on the syntax tree of t . If t is a variable, the claim is the bounded-part transfer inequality for F . If

$$t = \sigma(t_1, \dots, t_m),$$

then the induction hypothesis controls the costs of the transferred subterms

$$F(t_i(x_1, \dots, x_n)).$$

Operation compatibility identifies

$$F(\sigma_{\mathcal{X}}(t_1, \dots, t_m))$$

with

$$\sigma_{\mathcal{Y}}(F(t_1), \dots, F(t_m)).$$

The operation overhead α_{σ} then gives the recursive bound for t . $\square$

11 Base-Fibre Decompositions

Useful maps often lose information. Their value lies in making the lost information measurable.

Definition 11.1 (Relative presentation over a map). *Let*

$$q : \mathcal{X} \rightarrow \mathcal{Y}.$$

A relative presentation theory for the fibres of q assigns to each $y \in \mathcal{Y}$ a class $\mathcal{H}_y$ of relative descriptions, a relative realization map

$$\rho_y^{\text{rel}} : \mathcal{H}_y \rightarrow q^{-1}(y),$$

and a relative cost

$$\kappa_y^{\text{rel}} : \mathcal{H}_y \rightarrow B_{\text{rel}}.$$

The relative complexity is denoted

$$C_{\text{rel}}(x \mid q(x)).$$

Definition 11.2 (Controlled assembly). *The map q has controlled assembly with overhead A if there is an assembly procedure*

$$\text{Asm}(e, h) \in \mathcal{D}_{\mathcal{X}}$$

defined whenever e is a description of $y \in \mathcal{Y}$ and $h \in \mathcal{H}_y$, such that

$$\rho_{\mathcal{X}}(\text{Asm}(e, h)) = \rho_y^{\text{rel}}(h)$$

and

$$\kappa_{\mathcal{X}}(\text{Asm}(e, h)) \leq A(\kappa_{\mathcal{Y}}(e), \kappa_y^{\text{rel}}(h)).$$

Theorem 11.3 (Upper base-fibre bound). *If q has controlled assembly and either the relevant minimum costs are attained or A is compatible with the corresponding infima, then*

$$C_{\mathcal{X}}(x) \leq A(C_{\mathcal{Y}}(q(x)), C_{\text{rel}}(x \mid q(x))).$$

Proof. Choose a description e of $q(x)$ and a relative description h of x over $q(x)$. Applying the assembly procedure gives a description of x with cost at most

$$A(\kappa_{\mathcal{Y}}(e), \kappa_{q(x)}^{\text{rel}}(h)).$$

Taking infima over the two input descriptions gives the displayed bound, whenever the infima are attained or the overhead A is compatible with the relevant infima. In the usual discrete cost scales this is the same argument with minimum-cost descriptions. $\square$

Definition 11.4 (Controlled extraction). *The map q has controlled extraction if every description $d \in \mathcal{D}_{\mathcal{X}}$ of an object x yields:*

(i) *a description $\text{Base}(d)$ of $q(x)$ satisfying*

$$\kappa_{\mathcal{Y}}(\text{Base}(d)) \leq a(\kappa_{\mathcal{X}}(d));$$

(ii) *a relative description $\text{Rel}(d) \in \mathcal{H}_{q(x)}$ of x satisfying*

$$\kappa_{q(x)}^{\text{rel}}(\text{Rel}(d)) \leq r(\kappa_{\mathcal{X}}(d)).$$

Theorem 11.5 (Lower base-fibre bounds). *If q has controlled extraction and either the relevant minimum costs are attained or the overheads a and r are compatible with the corresponding infima, then*

$$C_Y(q(x)) \leq a(C_X(x))$$

and

$$C_{\text{rel}}(x \mid q(x)) \leq r(C_X(x)).$$

Consequently,

$$C_Y(q(x)) \not\leq a(b) \implies C_X(x) \not\leq b,$$

and

$$C_{\text{rel}}(x \mid q(x)) \not\leq r(b) \implies C_X(x) \not\leq b.$$

In totally ordered numerical scales the last two implications may be written with $>$.

Proof. In the minimum-cost case, apply extraction to a minimum-cost description of x . The extracted base and relative descriptions have costs controlled by a and r , giving the two upper bounds. The infimum version follows under the stated compatibility assumptions. The displayed implications are contrapositives. $\square$

Corollary 11.6 (Coarse base-fibre decomposition). *If both assembly and extraction are controlled in an overhead class $\mathcal{H}$, and the relevant minimum or infimum hypotheses above hold, then $C_X(x)$ is equivalent, up to $\mathcal{H}$ -distortion, to the joint data of $C_Y(q(x))$ and $C_{\text{rel}}(x \mid q(x))$.*

12 Verification Transfer

Proof data often move along reductions. The direction matters.

Definition 12.1 (Verification pullback). *Let $F : \mathcal{X} \rightarrow \mathcal{Y}$. A verification pullback from a target property P_Y to a source property P_X is a procedure*

$$\pi_Y \mapsto \pi_X$$

such that

$$R_Y(F(x), \pi_Y) \implies R_X(x, \pi_X)$$

and

$$\lambda_X(\pi_X) \leq h(\lambda_Y(\pi_Y)).$$

Definition 12.2 (Verification pushforward). *A verification pushforward is a procedure*

$$\pi_X \mapsto \pi_Y$$

such that

$$R_X(x, \pi_X) \implies R_Y(F(x), \pi_Y)$$

and

$$\lambda_Y(\pi_Y) \leq p(\lambda_X(\pi_X)).$$

Theorem 12.3 (Upper transfer for verification cost). *If a verification pullback exists with overhead h , then, in the extended resource scale and under the usual minimum or infimum-compatibility hypotheses,*

$$v_{P_X}(x) \leq h(v_{P_Y}(F(x))).$$

Proof. If $F(x)$ has no target verification data, then $v_{P_Y}(F(x)) = \infty$ and the inequality is vacuous after extending $h(\infty) = \infty$. Otherwise pull back arbitrary target verification data for $F(x)$ and use the overhead bound. Taking the infimum over target data gives the inequality whenever the infimum operation is compatible with h ; in discrete cost scales this follows by taking minimum-cost data. $\square$

Theorem 12.4 (Lower transfer for verification cost). *If a verification pushforward exists with overhead p , then, in the extended resource scale and under the usual minimum or infimum-compatibility hypotheses,*

$$v_{P_Y}(F(x)) \leq p(v_{P_X}(x)).$$

Hence

$$v_{P_Y}(F(x)) \not\leq p(s) \implies v_{P_X}(x) \not\leq s.$$

In totally ordered numerical scales the last implication may be written with $>$.

Proof. If x has no source verification data, then $v_{P_X}(x) = \infty$ and the first inequality is vacuous after extending $p(\infty) = \infty$. Otherwise push forward arbitrary source verification data and use the overhead bound. Taking infima gives the inequality whenever the infimum operation is compatible with p ; in discrete cost scales this follows from minimum-cost data. The lower-bound statement is the contrapositive. $\square$

13 Transfer Packages

Definition 13.1 (Transfer package). *Let $\mathfrak{P}_X$ and $\mathfrak{P}_Y$ be presentation contexts. A transfer package*

$$\mathfrak{T} : \mathfrak{P}_X \rightarrow \mathfrak{P}_Y$$

consists of some or all of the following compatible data:

- (i) an object map $F : \mathcal{X} \rightarrow \mathcal{Y}$;
- (ii) a description compiler $\hat{F} : \mathcal{D}_X \rightarrow \mathcal{D}_Y$;
- (iii) a presentation overhead f_{pres} ;
- (iv) an observable pullback $F^* : \Omega_Y \rightarrow \Omega_X$;
- (v) an observable overhead f_{obs} ;
- (vi) operation-compatibility data;
- (vii) verification-transfer data;
- (viii) fibre-control data;
- (ix) optional measure-distortion data.

The package is rich exactly on the layers for which these data are supplied.

Definition 13.2 (Overhead vector). *The overhead vector of a transfer package is*

$$\mathbf{f}_{\mathfrak{T}} = (f_{\text{pres}}, f_{\text{obs}}, f_{\text{op}}, f_{\text{ver}}, f_{\text{fib}}, f_{\text{meas}}),$$

with components interpreted only when the corresponding layer is present.

Theorem 13.3 (Composition of transfer packages). *If*

$$\mathfrak{T} : \mathfrak{P}_{\mathcal{X}} \rightarrow \mathfrak{P}_{\mathcal{Y}}$$

and

$$\mathfrak{S} : \mathfrak{P}_{\mathcal{Y}} \rightarrow \mathfrak{P}_{\mathcal{Z}}$$

are transfer packages, then their composition

$$\mathfrak{S} \circ \mathfrak{T} : \mathfrak{P}_{\mathcal{X}} \rightarrow \mathfrak{P}_{\mathcal{Z}}$$

is a transfer package on every layer controlled by both. The overheads compose componentwise.

Proof. The object maps and description compilers compose by the distortion theorem. Observable pullbacks compose contravariantly:

$$(F^* \circ G^*)(\theta) = F^*(\theta \circ G) = \theta \circ G \circ F.$$

Operation compatibility is obtained by pasting the operation diagrams. Verification-transfer maps compose in their declared directions. Fibre and measure controls compose by applying the corresponding bounds successively. $\square$

Remark 13.4 (Operational sheet for a transfer package). *For applications it is useful to record a transfer package in a fixed order:*

source context, target context, compiler,
target observables, observable pullback, verification data,
fibre map, moves in fibres, overhead vector,
statement to be transported.

The atlas of transfer packages is an expanded catalogue of such sheets. The theorem above explains why sheets can be composed; the controlled transfer schema below explains which logical statements are preserved by a completed sheet.

14 Budgeted Theorem Logic

To say that a transfer package transports a theorem, one must specify which statements are allowed.

Definition 14.1 (Budgeted fragment). *The budgeted fragment associated to a presentation context is generated by atoms of the following forms:*

$$\begin{aligned} x \in \mathcal{X}^{\leq b}, \quad \omega(x) \in S, \quad \omega(x) = \omega(x'), \\ R(x, \pi), \quad \lambda(\pi) \leq p, \quad x = t(x_1, \dots, x_n), \quad x \equiv x', \end{aligned}$$

using Boolean operations and bounded quantifiers such as

$$\forall x \in \mathcal{X}^{\leq b}, \quad \exists x \in \mathcal{X}^{\leq b}, \quad \forall \omega \in \Omega^{\leq r}, \quad \exists \pi : \lambda(\pi) \leq p.$$

Definition 14.2 (Interpretable formula). *A formula in the budgeted fragment of $\mathfrak{P}_{\mathcal{Y}}$ is interpretable by a transfer package*

$$\mathfrak{T} : \mathfrak{P}_{\mathcal{X}} \rightarrow \mathfrak{P}_{\mathcal{Y}}$$

if every atom and bounded quantifier appearing in the formula is supported by the corresponding layer of $\mathfrak{T}$.

For example, formulas involving bounded objects require presentation overhead; formulas involving observables require observable pullback; formulas involving verification data require verification transfer; target existential quantifiers require a controlled lifting or section; classification statements require reflection or fibre data.

The next statement is a theorem schema. Once a concrete budgeted language and its translation rules have been fixed, it becomes an ordinary induction on formulas.

Theorem 14.3 (Controlled transfer schema). *Let*

$$\mathfrak{T} : \mathfrak{P}_{\mathcal{X}} \rightarrow \mathfrak{P}_{\mathcal{Y}}$$

be a transfer package. Every universal theorem in the interpretable budgeted fragment of $\mathfrak{P}_{\mathcal{Y}}$ pulls back to a theorem in $\mathfrak{P}_{\mathcal{X}}$, with resource bounds transformed by the overhead vector of $\mathfrak{T}$.

In particular, if

$$\forall y \in \mathcal{Y}^{\leq c} \quad \Phi_{\mathcal{Y}}(y)$$

holds in the target and $f_{\text{pres}}(b) \leq c$, then

$$\forall x \in \mathcal{X}^{\leq b} \quad \Phi_{\mathcal{Y}}(F(x))$$

holds in the source. If additionally

$$\Phi_{\mathcal{Y}}(F(x)) \implies \Phi_{\mathcal{X}}(x)$$

is part of the transfer data, then

$$\forall x \in \mathcal{X}^{\leq b} \quad \Phi_{\mathcal{X}}(x).$$

Proof. The proof is by induction on formulas. For the atom $y \in \mathcal{Y}^{\leq c}$, bounded-part transfer gives

$$x \in \mathcal{X}^{\leq b}, \quad f_{\text{pres}}(b) \leq c \implies F(x) \in \mathcal{Y}^{\leq c}.$$

For observable atoms, replace a target observable η by the pulled-back observable $\eta \circ F$, with cost controlled by f_{obs} . For verification atoms, use the supplied verification pullback or pushforward, depending on the direction required by the statement. For operation atoms, use operation compatibility and term transfer. Equality or equivalence atoms transfer only when the package supplies the necessary reflection or fibre-control data.

Boolean connectives are immediate. Bounded universal quantifiers pull back because bounded source objects map into bounded target objects. Bounded existential quantifiers pull back only when the package includes controlled lifting data; this is precisely why interpretability is part of the hypothesis. The displayed universal statement is the special case with one bounded universal quantifier. $\square$

Remark 14.4 (No free transport). *The metatheorem does not say that every statement transfers. Existential target statements do not pull back without sections. Classifications do not transfer through lossy maps without fibre control. Verification-cost statements do not transfer without verification maps. Observable lower bounds do not transfer without observable compatibility.*

15 Fragmented Positive Budgets

Many non-computability results in presentation theory have the same formal shape. A positive property is semidecidable by searching through bounded fragments. Each bounded fragment is decidable, but the optimal bound needed on positive inputs is not computably controlled.

Definition 15.1 (Fragmented positive predicate). *Let $P \subseteq \Sigma^*$ be a semidecidable predicate. A decidable fragmentation of P is an increasing family*

$$P_{\leq 0} \subseteq P_{\leq 1} \subseteq P_{\leq 2} \subseteq \dots$$

such that

$$P = \bigcup_{r \in \mathbb{N}} P_{\leq r}$$

and membership in $P_{\leq r}$ is decidable uniformly in r .

For $x \in P$, define the positive width

$$w_P(x) = \min\{r : x \in P_{\leq r}\}.$$

The associated profile is

$$W_P(n) = \max\{w_P(x) : x \in P, |x| \leq n\},$$

with $W_P(n) = 0$ if the maximum is taken over the empty set.

Theorem 15.2 (Positive-fragment equivalence). *For every decidable fragmented positive predicate,*

$$W_P \equiv_T P.$$

In particular, if P is undecidable, then W_P has no computable majorant.

Proof. Assume first that an oracle for W_P is available. On input x , compute

$$N = W_P(|x|).$$

Because $P_{\leq N}$ is decidable, check whether

$$x \in P_{\leq N}.$$

If yes, then $x \in P$. If no, then $x \notin P$: indeed, if $x \in P$, then by definition of $W_P(|x|)$ every positive input of length at most $|x|$, including x , lies in $P_{\leq W_P(|x|)}$. Hence $P \leq_T W_P$.

Conversely, assume an oracle for P . To compute $W_P(n)$, enumerate the finite set of words $x \in \Sigma^*$ with $|x| \leq n$. Use the oracle to retain exactly those lying in P . For each retained word, search over $r = 0, 1, 2, \dots$ until the decidable test for $P_{\leq r}$ accepts. This search terminates because the word lies in P . Taking the maximum of the resulting finite list gives $W_P(n)$. Thus $W_P \leq_T P$.

If a computable function g satisfied $W_P(n) \leq g(n)$ for all n , the first reduction would decide P by replacing $W_P(|x|)$ with $g(|x|)$. Therefore an undecidable P admits no computable majorant for W_P . $\square$

Remark 15.3. *Finite-quotient search, bounded proof search, bounded normal-form search, and bounded move search are all instances of this template when the positive instances are exhausted by decidable bounded fragments.*

16 Fragmented Width Profiles and Effective Atlases

The previous theorem identifies the exact degree of the optimal positive width profile. Two refinements are often more useful in applications. First, any total upper bound for the width profile already computes the underlying positive predicate. Second, this gives a precise obstruction to effective atlases for access tasks.

Definition 16.1 (Majorant degrees). *For a fragmented positive predicate $P = \bigcup_r P_{\leq r}$, let*

$$\text{MajDeg}(W_P)$$

be the set of Turing degrees of total functions

$$G : \mathbb{N} \rightarrow \mathbb{N}$$

such that

$$W_P(n) \leq G(n) \quad \text{for every } n.$$

Theorem 16.2 (Majorant-cone theorem). *For every decidable fragmented positive predicate P ,*

$$\text{MajDeg}(W_P) = \{\mathbf{a} : \deg_T(P) \leq \mathbf{a}\}.$$

Proof. Let G be a total majorant of W_P . Given oracle access to G , decide P on input x by testing the decidable condition

$$x \in P_{\leq G(|x|)}.$$

If $x \in P$, then $w_P(x) \leq W_P(|x|) \leq G(|x|)$, so the test accepts. If the test accepts, soundness of the fragments gives $x \in P$. Hence every majorant computes P .

Conversely, if a degree $\mathbf{a}$ computes P , then it computes W_P by the positive-fragment equivalence. This gives at least one majorant,

$$G(n) = W_P(n)$$

of degree at most $\mathbf{a}$. To realize exactly $\mathbf{a}$, choose a total function g of degree $\mathbf{a}$. Since $W_P \leq_T g$, define

$$H(n) = 2^{W_P(n)}(2g(n) + 1).$$

Then H is a total majorant of W_P and $H \leq_T g$. Conversely $W_P(n)$ is the 2-adic valuation of $H(n)$, and then

$$g(n) = \frac{H(n)/2^{W_P(n)} - 1}{2}.$$

Thus $g \leq_T H$, so H has degree $\mathbf{a}$. The majorant degrees are precisely the upper cone above $\deg_T(P)$. $\square$

Definition 16.3 (Effective access atlas). *Let $P = \bigcup_r P_{\leq r}$ be a fragmented positive predicate. An effective access atlas for P is a computable system of local procedures, normal forms, coordinates, or search charts together with a computable total overhead function*

$$G : \mathbb{N} \rightarrow \mathbb{N}$$

such that every positive input $x \in P$ satisfies

$$x \in P_{\leq G(|x|)}.$$

Equivalently, the atlas supplies a computable total majorant for W_P .

Theorem 16.4 (No effective atlas criterion). *If P is undecidable, then P admits no effective access atlas. More generally, any oracle that computes the total overhead of such an atlas computes P .*

Proof. The overhead function of an effective access atlas is a total majorant of W_P . By the majorant-cone theorem, every such majorant computes P . Therefore a computable atlas overhead would decide P , contradicting undecidability. $\square$

Remark 16.5. *This is stronger than the absence of a computable classifier. It rules out any computable coordinate system that solves the declared access task with computable overhead: bounded observable refinement, bounded normal-form search, bounded fibre navigation, bounded proof degree, bounded finite quotient search, bounded model search, or bounded obstruction radius.*

Definition 16.6 (Cost-bounded reduction of fragmented predicates). *Let $P \subseteq X$ and $Q \subseteq Y$ be fragmented positive predicates on effectively encoded classes. A cost-bounded reduction from P to Q is a computable map*

$$\Phi : X \rightarrow Y$$

and a computable nondecreasing overhead

$$\alpha : \mathbb{N} \rightarrow \mathbb{N}$$

such that

$$x \in P \iff \Phi(x) \in Q$$

and

$$|\Phi(x)| \leq \alpha(|x|)$$

for all $x \in X$.

Definition 16.7 (Image width profile). *For a cost-bounded reduction $\Phi : P \rightarrow Q$, define*

$$W_Q^\Phi(n) = \max_{\substack{x \in P \\ |x| \leq n}} w_Q(\Phi(x)),$$

with value 0 if there are no positive inputs of size at most n .

Theorem 16.8 (Exact degree of image profiles). *Let $\Phi : P \rightarrow Q$ be a cost-bounded reduction between fragmented positive predicates. Then*

$$W_Q^\Phi \equiv_T P.$$

Consequently,

$$\text{MajDeg}(W_Q^\Phi) = \{\mathbf{a} : \deg_T(P) \leq \mathbf{a}\}.$$

Proof. Pull the fragmentation of Q back along Φ :

$$P_{\leq r}^\Phi = \{x : \Phi(x) \in Q_{\leq r}\}.$$

This is a decidable fragmentation of P , and its width profile is exactly W_Q^Φ . Apply the positive-fragment equivalence and the majorant-cone theorem. $\square$

Corollary 16.9 (Target profile inherits source hardness). *Let W_Q be the full width profile of Q . If P cost-bounded reduces to Q , then every total majorant of W_Q computes P .*

Proof. If H majorizes W_Q , then

$$G(n) = H(\alpha(n))$$

majorizes $W_Q^\Phi(n)$, because $|\Phi(x)| \leq \alpha(|x|)$. The preceding theorem gives $P \leq_T G \leq_T H$. $\square$

17 Immediate Budget Explosions

The preceding theorem is deliberately abstract. Its strength is that many classical undecidability theorems already come with decidable bounded fragments. The following examples record the quantitative profiles that are obtained without additional domain-specific work.

Theorem 17.1 (No computable finite-model budget). *Fix a finite relational signature containing at least one binary relation symbol. There is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that every first-order sentence φ over this signature with a finite model has a finite model of cardinality at most

$$f(|\varphi|).$$

Proof. For $r \in \mathbb{N}$, let $P_{\leq r}$ be the set of first-order sentences over the fixed signature having a model of size at most r . This set is decidable: there are only finitely many structures of size at most r over the fixed finite signature, and satisfaction of a first-order sentence in a finite structure is decidable.

The union of the $P_{\leq r}$ is the finite satisfiability problem for the signature. If a computable f bounded the size of the least finite model in terms of $|\varphi|$, then finite satisfiability would be decidable by checking all finite structures of size at most $f(|\varphi|)$. This contradicts Trakhtenbrot's theorem, which gives undecidability of finite satisfiability already for such signatures. $\square$

Definition 17.2 (Finite-model profile). *Let*

$$m(\varphi) = \min\{|M| : M \models \varphi, M \text{ finite}\}$$

when φ has a finite model, and define

$$M_{\text{fin}}(n) = \max\{m(\varphi) : |\varphi| \leq n, \varphi \text{ has a finite model}\}.$$

Corollary 17.3. *The profile M_{fin} has the same Turing degree as finite satisfiability over the fixed signature. In particular, M_{fin} is not computably majorized.*

Proof. This is the positive-fragment equivalence applied to the decidable fragments $P_{\leq r}$ above. $\square$

Theorem 17.4 (No computable Diophantine solution-height budget). *Fix an effective encoding of integer polynomials in finitely many variables. There is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that every polynomial equation

$$F(x_1, \dots, x_m) = 0, \quad F \in \mathbb{Z}[x_1, \dots, x_m],$$

with an integer solution has an integer solution $a \in \mathbb{Z}^m$ satisfying

$$\max_i |a_i| \leq f(|F|).$$

Proof. For $r \in \mathbb{N}$, let $D_{\leq r}$ be the set of encoded integer polynomials F for which there exists $a \in \mathbb{Z}^m$ with

$$F(a) = 0 \quad \text{and} \quad \max_i |a_i| \leq r.$$

The set $D_{\leq r}$ is decidable by finite search. The union of these fragments is the solvability problem for Diophantine equations over $\mathbb{Z}$, which is undecidable by the Davis–Putnam–Robinson–Matiyasevich theorem.

If a computable f as in the statement existed, Diophantine solvability over $\mathbb{Z}$ would be decidable by checking all integer tuples with $\max_i |a_i| \leq f(|F|)$. This contradiction proves the result. $\square$

Definition 17.5 (Diophantine height profile). *For solvable F , set*

$$h_{\mathbb{Z}}(F) = \min\{\max_i |a_i| : a \in \mathbb{Z}^m, F(a) = 0\},$$

and define

$$H_{\mathbb{Z}}(n) = \max\{h_{\mathbb{Z}}(F) : |F| \leq n, F = 0 \text{ is solvable over } \mathbb{Z}\}.$$

Corollary 17.6. *The profile $H_{\mathbb{Z}}$ has the same Turing degree as Diophantine solvability over $\mathbb{Z}$. In particular, $H_{\mathbb{Z}}$ is not computably majorized.*

Proof. Apply the positive-fragment equivalence to the fragments $D_{\leq r}$. $\square$

Theorem 17.7 (Proof-length explosion). *Let T be a recursively axiomatized formal theory whose theorem set*

$$\text{Thm}(T)$$

is undecidable. Fix an effective proof system for T . For a theorem φ , let

$$\ell_T(\varphi) = \min\{|\pi| : \pi \text{ is a formal proof of } \varphi \text{ in } T\}.$$

There is no computable function

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that

$$\ell_T(\varphi) \leq f(|\varphi|)$$

for every theorem φ of T .

Proof. For $r \in \mathbb{N}$, let $T_{\leq r}$ be the set of formulas having a T -proof of length at most r . This set is decidable by enumerating all strings of length at most r and checking which of them are valid formal proofs.

If a computable function f bounded $\ell_T(\varphi)$ in terms of $|\varphi|$, then theoremhood in T would be decidable: on input φ , enumerate all T -proofs of length at most $f(|\varphi|)$ and check whether one proves φ . This contradicts the assumed undecidability of $\text{Thm}(T)$. $\square$

Definition 17.8 (Proof-length profile). *Define*

$$L_T(n) = \max\{\ell_T(\varphi) : |\varphi| \leq n, \varphi \in \text{Thm}(T)\}.$$

Corollary 17.9. *The profile L_T has the same Turing degree as $\text{Thm}(T)$. In particular, if theoremhood in T is undecidable, then L_T is not computably majorized.*

Proof. This is the positive-fragment equivalence for the decidable fragments $T_{\leq r}$. $\square$

Theorem 17.10 (No computable mortality-word bound). *Fix an effectively presented class of finite matrix families over a computable field for which the matrix mortality problem is undecidable. There is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that every mortal family

$$\mathcal{A} = \{A_1, \dots, A_m\}$$

in the class has a zero product

$$A_{i_1} \cdots A_{i_\ell} = 0$$

with

$$\ell \leq f(|\mathcal{A}|).$$

Proof. For $r \in \mathbb{N}$, let $M_{\leq r}$ be the set of matrix families admitting a zero product of length at most r . This set is decidable by enumerating all words of length at most r in the generators and multiplying the corresponding matrices.

If a computable f as in the statement existed, then mortality would be decidable: on input $\mathcal{A}$, enumerate all products of length at most $f(|\mathcal{A}|)$, and check whether any of them is zero. The assumed bound makes the negative answer correct. This contradicts undecidability of mortality in the chosen class. $\square$

Theorem 17.11 (No computable nilpotency-time bound). *In any effective class of cellular automata for which nilpotency is undecidable, there is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that every nilpotent cellular automaton A in the class becomes uniformly nilpotent by time

$$f(|A|).$$

Proof. For fixed A and T , the assertion that A^T maps every configuration to the nilpotent uniform configuration is decidable: the value of a cell after T steps depends only on a finite window determined by T and the radius of the local rule, so one checks finitely many patterns.

If a computable bound f existed, nilpotency would be decidable by computing $T = f(|A|)$ and checking whether A^T is already uniformly nilpotent. This contradicts undecidability of nilpotency in the chosen class; for instance, Kari's theorem gives such undecidability for one-dimensional cellular automata. $\square$

Theorem 17.12 (No computable first-ambiguity bound). *There is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that every ambiguous context-free grammar G has a terminal word w with two distinct parse trees and

$$|w| \leq f(|G|).$$

Proof. For fixed G and N , it is decidable whether G has an ambiguous word of length at most N : enumerate all terminal words of length at most N , and for each word use a finite parsing procedure to decide whether it has at least two parse trees.

If a computable bound f existed, ambiguity of context-free grammars would be decidable by checking all terminal words of length at most $f(|G|)$. This contradicts the classical undecidability of the ambiguity problem for context-free grammars. $\square$

Theorem 17.13 (No computable null-area bound). *There is a finite group presentation P for which no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

has the following property: every word w of length at most n representing the identity in G_P admits a van Kampen diagram of area at most

$$f(n).$$

Proof. Choose a finitely presented group with undecidable word problem. If a computable function f as in the statement existed for some finite presentation P of this group, then the word problem would be decidable. Given a word w of length n , enumerate all van Kampen diagrams over P of area at most $f(n)$, and check whether any has boundary label w . If one is found, then $w = 1$ in G_P . If none is found, the assumed bound implies $w \neq 1$.

This contradicts the choice of G_P . Equivalently, the Dehn function of such a presentation is not computably majorized. $\square$

18 Finite-Window Obstruction Radius

The positive-fragment theorem bounds the first successful search in positive instances. A dual pattern occurs when global existence is equivalent to consistency on all finite windows, and failures are detected by a finite obstruction.

Definition 18.1 (Effective finite-window system). *An effective finite-window system consists of instances e , a global existence predicate $P(e)$, and decidable predicates*

$$L_N(e), \quad N \in \mathbb{N},$$

such that

$$P(e) \iff \forall N L_N(e).$$

For a negative instance, define its obstruction radius by

$$r(e) = \min\{N : L_N(e) \text{ fails}\}.$$

Theorem 18.2 (No computable obstruction radius). *Suppose P is undecidable in an effective finite-window system. Then there is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that every negative instance e satisfies

$$r(e) \leq f(|e|).$$

Proof. Assume such an f exists. Given e , compute $N = f(|e|)$ and decide the finite list of predicates

$$L_0(e), L_1(e), \dots, L_N(e).$$

If one fails, then $P(e)$ is false. If none fails, then $P(e)$ is true: otherwise e would be a negative instance with obstruction radius at most N , contradicting the search. This decides P , contrary to the hypothesis. $\square$

Theorem 18.3 (No computable finite-window bound for Wang tilings). *There is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that every finite Wang tileset T which does not tile the plane already fails to tile some square

$$[1, N]^2$$

with

$$N \leq f(|T|).$$

Proof. For a tileset T , let $L_N(T)$ be the decidable statement that the square $[1, N]^2$ admits a locally valid T -tiling. By compactness, T tiles the plane if and only if $L_N(T)$ holds for every N . Thus the obstruction radius theorem applies to the global predicate “ T tiles the plane”. Since the domino problem is undecidable by Berger’s theorem, no computable bound on the first failing square can exist. $\square$

19 Convergence Modulus Obstructions

Some access systems approximate a limiting quantity by finite budgets. If a threshold problem for the limit is undecidable, then the convergence to the limit cannot have a computable modulus.

Theorem 19.1 (No computable convergence modulus from undecidable thresholds). *Let $a_e \in \mathbb{R}$ be a family of real quantities with computable approximants*

$$a_{e,N}, \quad N \in \mathbb{N},$$

such that

$$a_e = \lim_{N \rightarrow \infty} a_{e,N}.$$

Assume there are rationals

$$\alpha < \beta$$

for which the promise problem

$$a_e \leq \alpha \quad \text{or} \quad a_e \geq \beta$$

is undecidable. Then there is no computable function

$$M(e, \varepsilon)$$

such that

$$|a_e - a_{e,N}| \leq \varepsilon$$

for every $N \geq M(e, \varepsilon)$.

Proof. Suppose such an M exists. Choose a rational

$$0 < \varepsilon < \frac{\beta - \alpha}{4}.$$

Given e , compute

$$N = M(e, \varepsilon).$$

Compute $a_{e,N}$ to error less than ε . In the case $a_e \leq \alpha$, the computed value is at most $\alpha + 2\varepsilon$, which is below $(\alpha + \beta)/2$. In the case $a_e \geq \beta$, the computed value is at least $\beta - 2\varepsilon$, which is above $(\alpha + \beta)/2$. Comparing with the midpoint therefore decides the promised threshold problem, contradiction. $\square$

Theorem 19.2 (No computable finite-dimensional strategy budget). *For the families of finite nonlocal games arising from the $\text{MIP}^* = \text{RE}$ undecidability theorem, there is no computable function*

$$D(G, \varepsilon)$$

such that

$$\omega_q(G) - \omega_{\leq D(G, \varepsilon)}(G) \leq \varepsilon$$

for every game G in the family and every rational $\varepsilon > 0$.

Proof. Let

$$\omega_{\leq d}(G)$$

denote the supremal winning probability over strategies of local dimension at most d . For fixed d , this is a finite-dimensional semialgebraic optimization problem and is effectively approximable. Moreover

$$\omega_q(G) = \sup_d \omega_{\leq d}(G).$$

If a computable dimension budget $D(G, \varepsilon)$ existed, then $\omega_q(G)$ would be computably approximable to arbitrary prescribed precision: compute $d = D(G, \varepsilon)$ and approximate the finite-dimensional value $\omega_{\leq d}(G)$. For a constant-gap promise

$$\omega_q(G) \leq \alpha \quad \text{or} \quad \omega_q(G) \geq \beta$$

with $\alpha < \beta$, choose $0 < \varepsilon < (\beta - \alpha)/4$, compute $d = D(G, \varepsilon)$, and approximate $\omega_{\leq d}(G)$ to error less than ε . Since

$$\omega_{\leq d}(G) \leq \omega_q(G)$$

and

$$\omega_q(G) - \omega_{\leq d}(G) \leq \varepsilon,$$

the same midpoint comparison as in the previous theorem decides the promised threshold. This contradicts the undecidability consequence of $\text{MIP}^* = \text{RE}$, which gives finite nonlocal games with undecidable constant-gap entangled-value thresholds. $\square$

20 Move-Fibre Recognition

The preceding theorem concerns positive search. A second recurring situation concerns a single fibre of a realization map. If the fibre is connected by elementary moves, the height needed to navigate back to a base point is exactly as hard as recognizing membership in that fibre.

Definition 20.1 (Computable move presentation). *A computable move presentation consists of:*

- (i) *a decidable class of descriptions $\mathcal{D} \subseteq \Sigma^*$;*
- (ii) *a realization map $\rho : \mathcal{D} \rightarrow \mathcal{X}$;*
- (iii) *a computable cost $\kappa : \mathcal{D} \rightarrow \mathbb{N}$, with finitely many descriptions of cost at most N , effectively enumerable for each N ;*
- (iv) *a symmetric computable elementary move relation $d \leftrightarrow d'$ preserving realization;*
- (v) *connected fibres: if $\rho(d) = \rho(e)$, then d and e are joined by a finite move path.*

Fix an object $x_* \in \mathcal{X}$ and a base description $d_* \in \rho^{-1}(x_*)$. Write

$$\text{Rec}_{x_*} = \{d \in \mathcal{D} : \rho(d) = x_*\}.$$

Definition 20.2 (Bridge height). *For $d \in \text{Rec}_{x_*}$, define*

$$h(d, d_*) = \min_{\gamma: d \rightsquigarrow d_*} \max_{e \in \gamma} \kappa(e),$$

where γ ranges over finite move paths from d to d_* . The bridge profile of the fibre is

$$B_{x_*}(n) = \max\left(\{\kappa(d_*)\} \cup \{h(d, d_*) : d \in \text{Rec}_{x_*}, \kappa(d) \leq n\}\right).$$

Theorem 20.3 (Move-fibre recognition equivalence). *For every computable move presentation as above,*

$$B_{x_*} \equiv_T \text{Rec}_{x_*}.$$

If Rec_{x_} is undecidable, then B_{x_*} has no computable majorant.*

Proof. First assume oracle access to B_{x_*} . Given $d \in \mathcal{D}$, set

$$N = \max\{B_{x_*}(\kappa(d)), \kappa(d)\}.$$

Enumerate the finite move graph consisting of descriptions reachable from d by paths whose vertices all have cost at most N . If d_* appears in this finite graph, then $d \in \text{Rec}_{x_*}$, because moves preserve realization.

If $d \in \text{Rec}_{x_*}$, then the definition of B_{x_*} gives a path from d to d_* of height at most N , so the search finds d_* . If the search does not find d_* , then $d \notin \text{Rec}_{x_*}$. Hence $\text{Rec}_{x_*} \leq_T B_{x_*}$.

Conversely, assume oracle access to Rec_{x_*} . To compute $B_{x_*}(n)$, enumerate all descriptions d with $\kappa(d) \leq n$, and use the oracle to keep exactly those in Rec_{x_*} . For each retained d , search over height bounds $H = 0, 1, 2, \dots$, enumerating the finite move graph of paths from d that stay within cost H , until d_* is reached. Connectedness of the fibre guarantees termination. The first successful H is $h(d, d_*)$. Taking the maximum over the finite retained list gives $B_{x_*}(n)$. Thus $B_{x_*} \leq_T \text{Rec}_{x_*}$.

If B_{x_*} had a computable majorant, the first reduction would decide Rec_{x_*} by searching below that computable bound. Therefore undecidable recognition implies no computable majorant. $\square$

21 Murphy Transfer

The algebraic counterpart of non-recursive budget transfer is universality transfer. A rank-profile compiler moves singularities from incidence geometry into fibres of an observable.

Definition 21.1 (Rank-profile compiler). *Let $\Omega : \mathcal{X} \rightarrow A$ be an observable on an algebraic moduli problem $\mathcal{X}$. A rank-profile compiler into $(\mathcal{X}, \Omega)$ assigns to each labelled subspace arrangement*

$$\mathbf{L} = (L_1, \dots, L_m) \subset V$$

an object

$$\mathcal{C}(\mathbf{L}) \in \mathcal{X}$$

so that the following conditions hold.

(i) **Rank-profile condition.** *The value $\Omega(\mathcal{C}(\mathbf{L}))$ depends only on the labelled rank function*

$$r_{\mathbf{L}}(S) = \dim \sum_{i \in S} L_i.$$

(ii) **Fibre faithfulness.** *On each locally closed stratum with fixed labelled rank function, $\mathcal{C}$ preserves the local moduli of arrangements up to simultaneous linear change of coordinates and smooth factors.*

(iii) **Algebraicity.** *The assignment $\mathbf{L} \mapsto \mathcal{C}(\mathbf{L})$ is algebraic on the relevant arrangement charts.*

Theorem 21.2 (Murphy transfer for rank profiles). *Suppose $(\mathcal{X}, \Omega)$ admits a rank-profile compiler. Then every finite-type singularity appearing in a labelled subspace-arrangement realization space appears, up to stable equivalence, in a fibre of Ω . In particular, by Mnëv–Sturmfels universality in the scheme-theoretic form of Lee–Vakil, every finite-type singularity over $\mathbb{Z}$ appears, up to stable equivalence, in a fibre of Ω .*

Proof. Fix a labelled rank function r . By the rank-profile condition, all compiled objects $\mathcal{C}(\mathbf{L})$ with

$$r_{\mathbf{L}} = r$$

lie in a single fibre of Ω . The algebraicity condition identifies the compiled family as an algebraic family inside that fibre. The fibre-faithfulness condition says that, locally on the fixed-rank arrangement stratum, this family has the same singularity type as the arrangement realization space, up to the smooth factors coming from coordinate choices and presentation parameters.

Thus every singularity appearing in a fixed-rank labelled arrangement realization space appears stably in an Ω -fibre. The Mnëv–Sturmfels–Lee–Vakil universality theorem supplies labelled incidence, matroid, or subspace-arrangement realization spaces with arbitrary finite-type singularities over $\mathbb{Z}$. Applying the compiler gives the asserted fibres. $\square$

22 Wild-Fibre Transfer

Murphy transfer concerns local algebraic singularities inside fibres. A complementary phenomenon concerns classification complexity inside a single fibre. The benchmark used here is simultaneous similarity of pairs of matrices:

$$(A, B) \sim (A', B')$$

if there exists an invertible matrix P such that

$$A' = PAP^{-1}, \quad B' = PBP^{-1}.$$

Definition 22.1 (Matrix-pair subproblem in fixed fibres). *Let $\Omega : \mathcal{X} \rightarrow A$ be an observable. A matrix-pair subproblem in fixed fibres consists of maps*

$$\mathcal{C}_n : \mathcal{U}_n \rightarrow \mathcal{X}, \quad \mathcal{U}_n \subseteq M_n(k)^2,$$

and values $\omega_n \in A$, for infinitely many n , such that

$$\Omega(\mathcal{C}_n(A, B)) = \omega_n$$

for all $(A, B) \in \mathcal{U}_n$, and

$$\mathcal{C}_n(A, B) \cong \mathcal{C}_n(A', B') \iff (A, B) \sim (A', B')$$

within the chosen matrix-pair class.

If all ω_n are equal to a single value ω_0 , the subproblem lies inside one fibre.

Theorem 22.2 (Wild-fibre transfer). *If fixed fibres of an observable contain a matrix-pair subproblem whose source class contains simultaneous similarity of arbitrary matrix pairs by a fully faithful representation embedding, then classification within those fixed fibres is wild. If all ω_n are equal, the same conclusion holds inside one fibre.*

Proof. For each n , the map $\mathcal{C}_n$ sends all source objects to the fixed fibre $\Omega^{-1}(\omega_n)$. The displayed equivalence says that isomorphism inside that fibre, restricted to this image, is exactly simultaneous similarity of the source pairs. Therefore any classification of objects in these fibres would, by restriction to the image of the $\mathcal{C}_n$, classify the source matrix-pair problem. If the source contains arbitrary simultaneous similarity by a fully faithful representation embedding, then the corresponding fixed fibres contain a standard wild classification problem. When all ω_n coincide, this subproblem lies inside the single fibre $\Omega^{-1}(\omega_0)$. $\square$

23 Budgeted Morita Equivalence

Transfer packages can also express when two presentation contexts are quantitatively the same theory.

Definition 23.1 (Budgeted Morita equivalence). *Let $\mathcal{H}$ be an overhead class. Two presentation contexts $\mathfrak{P}$ and $\mathfrak{Q}$ are $\mathcal{H}$ -Morita equivalent if there are transfer packages*

$$\mathfrak{P} \longrightarrow \mathfrak{Q}, \quad \mathfrak{Q} \longrightarrow \mathfrak{P},$$

with all overheads in $\mathcal{H}$, such that the two composites are equivalent to the identity packages on objects, bounded parts, observables, operations, verification relations, and fibre navigation profiles, again with overheads in $\mathcal{H}$.

Proposition 23.2 (Invariance under budgeted Morita equivalence). *If $\mathfrak{P}$ and $\mathfrak{Q}$ are $\mathcal{H}$ -Morita equivalent, then every theorem in the interpretable budgeted fragment of one context transfers to the other with $\mathcal{H}$ -distortion. In particular, presentation complexity, observable distinguishing costs, verification profiles, bridge profiles, and bounded obstruction profiles agree up to $\mathcal{H}$ -equivalence whenever they are expressible in that fragment.*

Proof. Apply the controlled transfer schema to the package $\mathfrak{P} \rightarrow \mathfrak{Q}$ and then to the package $\mathfrak{Q} \rightarrow \mathfrak{P}$. The comparison data for the composites identify the transported statements with the original ones, while the closure of $\mathcal{H}$ under composition keeps the resulting overheads inside $\mathcal{H}$. The listed profiles are defined by bounded existential or universal statements over the layers explicitly preserved by the equivalence, so their upper and lower bounds transport in both directions. $\square$

24 Transfer Graphs

Definition 24.1 (Transfer graph). *A transfer graph is a directed graph whose vertices are presentation contexts and whose arrows are transfer packages. Each arrow is labelled by its overhead vector.*

Definition 24.2 (Transport along a path). *A path*

$$\mathfrak{P}_0 \rightarrow \mathfrak{P}_1 \rightarrow \cdots \rightarrow \mathfrak{P}_n$$

induces a composite transfer package

$$\mathfrak{P}_0 \rightarrow \mathfrak{P}_n.$$

Its overhead vector is obtained by componentwise composition of the overheads along the path.

The resulting transport distance is generally not a metric. It is directed, multiresource, and takes values in functions or upper sets of functions.

Common hubs in transfer graphs include representation schemes, matrix-pair classification, finite-window constraint systems, computable limiting processes, nonlocal games, cellular automata, formal grammars, matrix semigroups, chain complexes, tensor networks, tropicalizations, finite-test profiles, proof systems, automata, syntactic monoids, differential modules, and Galois representations.

Remark 24.3 (Transfer classes and moduli). *Most arguments in this note use a single transfer package at a time. In some problems, however, the useful object is the class of all packages between two contexts with a prescribed overhead profile. One may write informally*

$$\text{Trans}_{\leq \chi}(\mathfrak{P}, \mathfrak{Q})$$

for transfer packages from $\mathfrak{P}$ to $\mathfrak{Q}$ whose cost overhead, observable loss, verification overhead, and fibre distortion are bounded by a character χ . This notation is only a bookkeeping device here, not a new layer of foundations. It becomes useful when local transfer packages must be compared or glued, when one wants to optimize over several possible proof strategies, or when a capacity argument rules out every package in a proposed class. In that sense, a transfer package proves a transported theorem, while a transfer class records a family of possible proof strategies.

25 Application I: Rank-Invariant Fibres

We first apply the transfer viewpoint to multiparameter persistence. The point is not that the rank invariant is incomplete; that is well known. The point is that its fibres can contain arbitrary algebraic singularity types.

Let k be an algebraically closed field and let

$$R = k[t_1, \dots, t_m]$$

with its standard $\mathbb{N}^m$ -grading. A finitely presented $\mathbb{N}^m$ -graded R -module is an m -parameter persistence module.

For such a module M , the rank invariant is

$$\rho_M(a, b) = \text{rank}(M_a \rightarrow M_b), \quad a \leq b.$$

25.1 The Incidence Compiler

Let V be a finite-dimensional k -vector space and let

$$\mathbf{L} = (L_1, \dots, L_m)$$

be a labelled tuple of subspaces $L_i \subseteq V$. Define

$$P(\mathbf{L}) = (R \otimes_k V) \Big/ \sum_{i=1}^m t_i R \otimes_k L_i.$$

This module is generated in degree 0, with relations $t_i \ell = 0$ for $\ell \in L_i$.

Lemma 25.1 (Graded pieces). *For $a = (a_1, \dots, a_m) \in \mathbb{N}^m$, set*

$$S(a) = \{i : a_i \geq 1\}.$$

Then

$$P(\mathbf{L})_a \simeq V \Big/ \sum_{i \in S(a)} L_i.$$

If $a \leq b$, the structure map

$$P(\mathbf{L})_a \rightarrow P(\mathbf{L})_b$$

is the natural quotient map

$$V \Big/ \sum_{i \in S(a)} L_i \longrightarrow V \Big/ \sum_{i \in S(b)} L_i.$$

Proof. The degree- a component of $R \otimes_k V$ is a copy of V , generated by $t^a \otimes V$. The submodule $t_i R \otimes L_i$ contributes to degree a exactly when $a_i \geq 1$, and in that degree it contributes the subspace $L_i \subseteq V$. Hence the degree- a quotient is

$$V \Big/ \sum_{i: a_i \geq 1} L_i.$$

Multiplication by t^{b-a} carries the degree- a copy of V to the degree- b copy of V , and the relation subspace can only increase from $\sum_{i \in S(a)} L_i$ to $\sum_{i \in S(b)} L_i$. Therefore the induced map is the displayed quotient map. $\square$

Lemma 25.2 (Rank invariant as a subspace-rank function). *For $a \leq b$,*

$$\rho_{P(\mathbf{L})}(a, b) = \dim V - \dim \left(\sum_{i \in S(b)} L_i \right).$$

Thus the rank invariant of $P(\mathbf{L})$ depends only on the labelled rank function

$$r_{\mathbf{L}}(S) = \dim \sum_{i \in S} L_i.$$

Proof. By the preceding lemma, $P(\mathbf{L})_a \rightarrow P(\mathbf{L})_b$ is the quotient map

$$V / \sum_{i \in S(a)} L_i \longrightarrow V / \sum_{i \in S(b)} L_i.$$

This map is surjective. Its rank is therefore the dimension of its codomain:

$$\dim V - \dim \sum_{i \in S(b)} L_i.$$

$\square$

Lemma 25.3 (Recovery of the labelled arrangement). *Let $\mathbf{L} = (L_1, \dots, L_m) \subset V$ and $\mathbf{L}' = (L'_1, \dots, L'_m) \subset V'$. Then*

$$P(\mathbf{L}) \cong P(\mathbf{L}')$$

as $\mathbb{N}^m$ -graded R -modules if and only if there is a linear isomorphism

$$T : V \rightarrow V'$$

such that

$$T(L_i) = L'_i$$

for every i .

Proof. A graded R -module homomorphism

$$\Phi : P(\mathbf{L}) \rightarrow P(\mathbf{L}')$$

is determined by its degree-zero part $T : V \rightarrow V'$, because $P(\mathbf{L})$ is generated in degree 0. In $P(\mathbf{L})$, the subspace $L_i \subset V = P(\mathbf{L})_0$ is exactly

$$\ker(P(\mathbf{L})_0 \xrightarrow{t_i} P(\mathbf{L})_{e_i}),$$

since this map is $V \rightarrow V/L_i$.

Because Φ commutes with multiplication by t_i , the map T sends this kernel into the corresponding kernel L'_i . Thus

$$T(L_i) \subseteq L'_i.$$

If Φ is an isomorphism, applying the same argument to Φ^{-1} gives equality.

Conversely, if $T : V \rightarrow V'$ is an isomorphism satisfying $T(L_i) = L'_i$ for all i , then

$$\text{id}_R \otimes T : R \otimes V \rightarrow R \otimes V'$$

sends $t_i R \otimes L_i$ onto $t_i R \otimes L'_i$. It therefore descends to an isomorphism

$$P(\mathbf{L}) \cong P(\mathbf{L}').$$

$\square$

Theorem 25.4 (Universality inside rank-invariant fibres). *In the standard scheme-theoretic presentation of finite multigraded R -modules of the form above, rank-invariant fibres contain locally closed strata stably equivalent to realization spaces of labelled subspace arrangements. Consequently, by Mnëv–Sturmfels universality in scheme-theoretic form, every finite-type singularity over $\mathbb{Z}$ occurs, up to stable equivalence, in a rank-invariant fibre of finitely presented multiparameter persistence modules.*

Proof. The construction $\mathbf{L} \mapsto P(\mathbf{L})$ is algebraic in the Pluecker coordinates of the labelled subspaces. The preceding lemmas show two facts.

First, the rank invariant of $P(\mathbf{L})$ is determined by the rank function

$$S \mapsto \dim \sum_{i \in S} L_i.$$

Therefore all labelled arrangements with the same rank function map into a single rank-invariant fibre.

Second, within the image of this construction, the graded module $P(\mathbf{L})$ remembers the labelled arrangement up to simultaneous linear change of coordinates. Thus the corresponding locally closed piece of the rank-invariant fibre has the same local moduli, up to the usual stable factors coming from choices of coordinates and presentation format, as the realization space of the arrangement.

Scheme-theoretic Mnëv–Sturmfels universality says that realization spaces of finite incidence or matroid-type configurations realize every finite-type singularity over $\mathbb{Z}$, up to stable equivalence. Such incidence configurations can be encoded as labelled subspace arrangements: for example, points and lines in $\mathbb{P}^2$ become one- and two-dimensional subspaces of k^3 , and incidence is the rank condition

$$p \subseteq \ell \iff \dim(p + \ell) = 2.$$

Non-incidence is the complementary locally open rank condition. Hence the relevant realization spaces appear among fixed-rank-function subspace-arrangement strata. Applying the compiler $\mathbf{L} \mapsto P(\mathbf{L})$ places these strata inside rank-invariant fibres. $\square$

Corollary 25.5. *Rank-invariant fibres in multiparameter persistence can be reducible, non-normal, non-reduced, and singular in arbitrary finite-type ways.*

Theorem 25.6 (Restricted real rank-profile realizability). *Over $\mathbb{R}$, the following restricted realizability problem is $\exists\mathbb{R}$ -complete: given integers d, m and a function*

$$r : 2^{\{1, \dots, m\}} \rightarrow \mathbb{N},$$

decide whether there are subspaces

$$L_1, \dots, L_m \subseteq \mathbb{R}^d$$

such that

$$\dim \sum_{i \in S} L_i = r(S)$$

for every $S \subseteq \{1, \dots, m\}$. Consequently, the rank-invariant realizability problem for persistence modules in the image of the compiler $\mathbf{L} \mapsto P(\mathbf{L})$ is $\exists\mathbb{R}$ -complete.

Proof. Membership in $\exists\mathbb{R}$ follows by choosing matrix coordinates for the subspaces L_i . The condition

$$\dim \sum_{i \in S} L_i \leq r(S)$$

is expressed by vanishing of all $(r(S) + 1)$ -minors of the concatenated matrix for the L_i with $i \in S$. The condition

$$\dim \sum_{i \in S} L_i \geq r(S)$$

is expressed by the non-vanishing of at least one $r(S)$ -minor, equivalently by a finite disjunction of polynomial inequalities. Since existential first-order formulas over the reals allow finite Boolean combinations of polynomial equalities and inequalities, this gives an $\exists\mathbb{R}$ description.

Hardness follows from real representability of rank-three matroids. A rank-three matroid on labelled elements can be encoded by one-dimensional subspaces

$$L_i \subseteq \mathbb{R}^3,$$

with the matroid rank of S equal to

$$\dim \sum_{i \in S} L_i.$$

Mnëv universality, in particular its standard complexity-theoretic consequence for real matroid realizability, gives $\exists\mathbb{R}$ -hardness. The final statement follows from the lemma identifying the rank invariant of $P(\mathbf{L})$ with the labelled rank function of $\mathbf{L}$. $\square$

Theorem 25.7 (Boolean-lattice restriction-rank fibres). *Let B_m be the Boolean lattice of subsets of $\{1, \dots, m\}$. For a labelled subspace arrangement*

$$\mathbf{L} = (L_1, \dots, L_m) \subset V,$$

define a functor

$$F_{\mathbf{L}} : B_m \rightarrow \text{Vect}_k$$

by

$$F_{\mathbf{L}}(S) = V / \sum_{i \in S} L_i,$$

with the natural quotient map $F_{\mathbf{L}}(S) \rightarrow F_{\mathbf{L}}(T)$ for $S \subseteq T$. The restriction-rank profile

$$\rho_{\mathbf{L}}(S, T) = \text{rank}(F_{\mathbf{L}}(S) \rightarrow F_{\mathbf{L}}(T))$$

has fibres containing, up to stable equivalence, every finite-type singularity over $\mathbb{Z}$.

Proof. For $S \subseteq T$, the map

$$F_{\mathbf{L}}(S) \rightarrow F_{\mathbf{L}}(T)$$

is surjective. Hence

$$\rho_{\mathbf{L}}(S, T) = \dim F_{\mathbf{L}}(T) = \dim V - \dim \sum_{i \in T} L_i.$$

Thus the full restriction-rank profile depends only on the labelled rank function

$$T \mapsto \dim \sum_{i \in T} L_i.$$

Conversely, the functor $F_{\mathbf{L}}$ remembers the labelled arrangement up to simultaneous linear change of coordinates, because

$$F_{\mathbf{L}}(\emptyset) = V$$

and

$$L_i = \ker(F_{\mathbf{L}}(\emptyset) \rightarrow F_{\mathbf{L}}(\{i\})).$$

Therefore the map $\mathbf{L} \mapsto F_{\mathbf{L}}$ is a rank-profile compiler in the sense of the Murphy transfer theorem. Applying that theorem gives the asserted singularity types inside fibres of the restriction-rank profile. $\square$

Theorem 25.8 (Fixed rank-profile fibres contain matrix-pair classification). *Let k be an algebraically closed field. For each n , there is a restriction-rank profile on the Boolean lattice B_5 whose fibre contains the simultaneous-similarity classification of an open matrix-pair class in $M_n(k)^2$. As n varies, these fixed-profile fibres contain a fully faithful copy of the classification of arbitrary pairs of matrices up to simultaneous similarity.*

Proof. Let $U = k^n$ and $W = U \oplus U$. For a pair $(A, B) \in M_n(k)^2$, define five labelled n -dimensional subspaces of W :

$$\begin{aligned} E &= U \oplus 0, & F &= 0 \oplus U, & D &= \{(x, x) : x \in U\}, \\ G_A &= \{(x, Ax) : x \in U\}, & G_B &= \{(x, Bx) : x \in U\}. \end{aligned}$$

Assume

$$A, B, A - I, B - I, A - B$$

are invertible. Then any two distinct subspaces in the list span W . Therefore the labelled rank function of this five-subspace arrangement is independent of (A, B) :

$$\dim \sum_{i \in S} L_i = \begin{cases} 0, & S = \emptyset, \\ n, & |S| = 1, \\ 2n, & |S| \geq 2. \end{cases}$$

Applying the Boolean-lattice construction to this arrangement gives functors

$$F_{A,B} : B_5 \rightarrow \text{Vect}_k$$

all lying in one restriction-rank fibre for the fixed value above.

It remains to identify isomorphism in this fibre on the constructed subfamily. The functor remembers the labelled arrangement because each labelled subspace is recovered as

$$L_i = \ker(F_{A,B}(\emptyset) \rightarrow F_{A,B}(\{i\})).$$

Thus an isomorphism $F_{A,B} \cong F_{A',B'}$ is the same as a linear isomorphism $T : W \rightarrow W$ preserving the five labelled subspaces. Preservation of E and F gives

$$T = \begin{pmatrix} P & 0 \\ 0 & Q \end{pmatrix}.$$

Preservation of D forces $Q = P$. Preservation of G_A and G_B is then exactly

$$A' = PAP^{-1}, \quad B' = PBP^{-1}.$$

The converse is immediate from the block-diagonal map $\text{diag}(P, P)$. Hence the fixed rank-profile fibre contains simultaneous similarity of the indicated open matrix-pair class.

The open conditions do not remove the usual matrix-pair classification problem. Given an arbitrary pair $(X, Y) \in M_n(k)^2$, choose scalars $\lambda_i, \mu_i \in k \setminus \{0, 1\}$ with the λ_i distinct and $\lambda_i \neq \mu_i$. Define, on $U^{\oplus 5}$,

$$A_0 = \text{diag}(\lambda_1 I, \dots, \lambda_5 I)$$

and

$$B_{X,Y} = \begin{pmatrix} \mu_1 I & I & X & Y & 0 \\ 0 & \mu_2 I & I & 0 & 0 \\ 0 & 0 & \mu_3 I & I & 0 \\ 0 & 0 & 0 & \mu_4 I & I \\ 0 & 0 & 0 & 0 & \mu_5 I \end{pmatrix}.$$

The pair $(A_0, B_{X,Y})$ satisfies the invertibility conditions above. If

$$(A_0, B_{X,Y}) \sim (A_0, B_{X',Y'}),$$

then the conjugating matrix must be block diagonal because the λ_i are distinct. The identity blocks in $B_{X,Y}$ force all diagonal blocks to be equal, and the $(1,3)$ - and $(1,4)$ -blocks then give

$$X' = PXP^{-1}, \quad Y' = PYP^{-1}.$$

Conversely, any simultaneous similarity between (X, Y) and (X', Y') gives a simultaneous similarity between the stabilized pairs. Thus arbitrary matrix-pair similarity embeds fully faithfully into the open class used above. $\square$

Corollary 25.9 (Fixed rank-invariant fibres in five-parameter persistence). *For each n , there is a rank invariant ρ_n of finitely presented 5-parameter persistence modules such that the fibre*

$$\{M : \rho_M = \rho_n\}$$

contains the simultaneous-similarity classification of an open class of pairs in $M_n(k)^2$. As n varies, these fixed-rank-invariant fibres contain the usual matrix-pair classification problem.

Proof. Use the same five subspaces

$$(E, F, D, G_A, G_B) \subset W = U \oplus U$$

as in the preceding theorem, and form the 5-parameter module

$$M_{A,B} = (k[t_1, \dots, t_5] \otimes_k W) \Big/ \sum_{i=1}^5 t_i k[t_1, \dots, t_5] \otimes_k L_i.$$

The graded-piece lemma gives

$$(M_{A,B})_a \simeq W \Big/ \sum_{i \in S(a)} L_i, \quad S(a) = \{i : a_i \geq 1\}.$$

Therefore the rank invariant is determined by

$$S \mapsto \dim \sum_{i \in S} L_i.$$

Under the open conditions

$$A, B, A - I, B - I, A - B$$

invertible, this rank function is the fixed function

$$\dim \sum_{i \in S} L_i = \begin{cases} 0, & S = \emptyset, \\ n, & |S| = 1, \\ 2n, & |S| \geq 2. \end{cases}$$

Hence all $M_{A,B}$ lie in one rank-invariant fibre.

The recovery lemma for $P(\mathbf{L})$ says that an isomorphism

$$M_{A,B} \cong M_{A',B'}$$

is exactly a simultaneous linear isomorphism of the five labelled subspaces. The proof of the preceding theorem identifies this condition with

$$A' = PAP^{-1}, \quad B' = PBP^{-1}.$$

The same stabilization argument embeds arbitrary matrix-pair similarity into the open class. $\square$

Proposition 25.10 (Finite-poset presentation equivalence). *For every finite poset P , the following presentation contexts are polynomially equivalent:*

functors $P \rightarrow \text{Vect}_k$,

representations of the Hasse quiver of P with commutativity relations,

finite modules over the incidence algebra kP ,

and sheaves or cosheaves on the Alexandrov space associated to P , with the corresponding choice of covariance. Under these translations, dimension vectors, ranks of structure maps, morphisms, endomorphism algebras, and decomposability are preserved with polynomial overhead in the size of P and the chosen linear data.

Proof. A functor $P \rightarrow \text{Vect}_k$ assigns a vector space to each $p \in P$ and a linear map to each relation $p \leq q$, compatible with composition. The Hasse quiver records only cover relations. Imposing commutativity along all comparable paths reconstructs exactly the functorial data.

The incidence algebra kP has idempotents e_p and basis elements e_{pq} for $p \leq q$. A finite kP -module decomposes into components $e_p M$, and the elements e_{pq} give compatible maps $e_p M \rightarrow e_q M$. This is again the same data as a functor $P \rightarrow \text{Vect}_k$.

For the Alexandrov topology associated to P , sheaves and cosheaves encode the same finite diagram, with variance determined by the convention for specialization order. All conversions are explicit and use only finitely many linear maps indexed by P , so their presentation overhead is polynomial. $\square$

Corollary 25.11 (Rank-derived invariants inherit fixed-fibre complexity). *Let I be any invariant on one of the finite-poset, sheaf, quiver, or multiparameter-persistence contexts above that factors through the relevant rank profile:*

$$I(M) = \Phi(\rho_M).$$

Then some fibre of I contains, up to stable equivalence, every finite-type singularity over $\mathbb{Z}$. Moreover, for each n , some fibre of I contains the simultaneous-similarity classification of pairs of $n \times n$ matrices. Thus rank-derived invariants do not remove the fixed-profile matrix-pair subproblems exhibited above.

Proof. By the Boolean-lattice theorem, for every finite-type singularity there is a rank-profile fibre containing that singularity stably. If $I = \Phi \circ \rho$, then every rank-profile fibre is contained in a fibre of I . Therefore the same singularity occurs in an I -fibre.

For the matrix-pair statement, use either construction above. The objects constructed there have a fixed rank profile ρ_n , hence a fixed value

$$\Phi(\rho_n)$$

of I . Their isomorphism problem is simultaneous similarity of $n \times n$ matrix pairs. $\square$

26 Application II: Finite-Quotient Observable Budgets

Finite quotients form a natural observable system for finitely presented groups.

Let $P = \langle S \mid R \rangle$ be a finite presentation, and let G_P be the group it presents. Let $|P|$ denote a fixed total presentation length.

For every finite group Q , define the observable

$$\omega_Q(G) = \text{Hom}(G, Q),$$

with observable cost

$$\chi(\omega_Q) = |Q|.$$

Definition 26.1 (First nontrivial finite quotient). *Define*

$$q(P) = \min\{|Q| : Q \text{ finite and there exists } G_P \rightarrow Q \text{ with nontrivial image}\}.$$

If no such Q exists, set $q(P) = \infty$.

Theorem 26.2 (No computable finite-quotient observable budget). *There is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that, for every finite presentation P ,

$$q(P) < \infty \implies q(P) \leq f(|P|).$$

Proof. Assume such an f exists. Given a finite presentation P , compute

$$N = f(|P|).$$

There are finitely many multiplication tables of groups of order at most N , and they can be enumerated. For each such finite group Q , enumerate all maps from the finite generating set S of P to Q . Each map determines a homomorphism from the free group on S . Check whether every relator in R maps to the identity of Q . If so, the map descends to a homomorphism

$$G_P \rightarrow Q.$$

Then check whether its image is nontrivial.

If such a homomorphism is found, G_P has a nontrivial finite quotient. If no such homomorphism is found for any Q of order at most N , the assumed bound implies that no nontrivial finite quotient exists. Thus f would decide whether the profinite completion of G_P is nontrivial.

This contradicts the undecidability theorem of Bridson and Wilton for triviality of profinite completions of finitely presented groups. Therefore no computable f exists. $\square$

Corollary 26.3 (Invisibility below every computable finite budget). *For every computable function f , there exists a finite presentation P such that G_P has a nontrivial finite quotient, but every homomorphism*

$$G_P \rightarrow Q$$

is trivial for every finite group Q with

$$|Q| \leq f(|P|).$$

Proof. If the statement failed for some computable f , then f would bound $q(P)$ for every P with $q(P) < \infty$, contradicting the theorem. $\square$

Definition 26.4 (Finite-quotient profile). *Define*

$$W(n) = \max\{q(P) : |P| \leq n, q(P) < \infty\}.$$

Proposition 26.5 (Turing degree of the finite-quotient profile). *The function W has the same Turing degree as the decision problem*

$$A = \{P : G_P \text{ has a nontrivial finite quotient}\}.$$

Proof. Given oracle access to W , decide A as follows. On input P , compute

$$N = W(|P|).$$

Enumerate finite groups Q of order at most N , and enumerate homomorphisms $G_P \rightarrow Q$. If a nontrivial image is found, answer yes; otherwise answer no. By definition of W , this is correct. Hence $A \leq_T W$.

Conversely, given oracle access to A , compute $W(n)$ by enumerating all finite presentations P with $|P| \leq n$. There are finitely many. Use the oracle to select those in A . For each selected P , enumerate finite groups and homomorphisms until the least nontrivial finite quotient is found. This search terminates for every selected P . Take the maximum of the resulting finite list. Hence $W \leq_T A$. $\square$

Corollary 26.6 (No computable finite-linear observable budget). *There is no computable function*

$$d : \mathbb{N} \rightarrow \mathbb{N}$$

such that, for every finite presentation P , if G_P has a nontrivial finite quotient, then there exists a nontrivial homomorphism

$$G_P \rightarrow \mathrm{GL}_{d'}(\mathbb{F}_2)$$

for some $d' \leq d(|P|)$.

Proof. If such a computable function existed, then one could decide whether G_P has a nontrivial finite quotient as follows. Compute $d(|P|)$. For each

$$1 \leq d' \leq d(|P|),$$

the finite group $\mathrm{GL}_{d'}(\mathbb{F}_2)$ is explicitly enumerable. Enumerate all maps from the generators of P to this group, check the relators, and test whether the resulting image is nontrivial. If a nontrivial image is found, answer yes; otherwise answer no. The assumed bound makes the negative answer correct.

This contradicts Bridson–Wilton undecidability. Conversely, every nontrivial finite quotient Q has a faithful left-regular permutation representation, hence a faithful linear representation

$$Q \hookrightarrow \mathrm{GL}_{|Q|}(\mathbb{F}_2).$$

Thus finite-linear observables over $\mathbb{F}_2$ detect exactly the same positive property, but they do not admit a computable completeness budget. $\square$

27 Application III: Pachner Fibre Barriers

The third application concerns the geometry of a single fibre of a triangulation presentation.

Let $\mathcal{D}$ be the class of finite triangulations of closed PL 4-manifolds. The realization map is

$$\rho(T) = |T|_{\mathrm{PL}},$$

and the cost is

$$\kappa(T) = \#\{4\text{-simplices of } T\}.$$

For a fixed closed PL 4-manifold M , the fibre

$$\mathcal{F}(M) = \rho^{-1}(M)$$

is the set of all triangulations of M . Pachner moves connect this fibre.

Definition 27.1 (Pachner height). *Fix a triangulation $T_* \in \mathcal{F}(M)$. For $T \in \mathcal{F}(M)$, define*

$$h_P(T, T_*) = \min_{\gamma: T \rightsquigarrow T_*} \max_{S \in \gamma} |S|,$$

where γ ranges over Pachner paths from T to T_* , and $|S|$ denotes the number of 4-simplices of S .

Theorem 27.2 (Non-recursive Pachner bridge height). *There exists a closed PL 4-manifold M and a triangulation T_* of M such that no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

satisfies

$$h_P(T, T_*) \leq f(|T|)$$

for every triangulation T of M .

Proof. Choose a closed PL 4-manifold M that is not algorithmically recognizable: there is no algorithm deciding, for an arbitrary triangulated closed PL 4-manifold X , whether $X \cong_{\text{PL}} M$. Such manifolds exist by Markov-type unrecognizability results in dimension 4.

Fix a triangulation T_* of M . Suppose, toward contradiction, that a computable function f bounds $h_P(T, T_*)$ for every triangulation T of M .

We construct an algorithm recognizing M . Given a triangulated closed PL 4-manifold X , compute

$$N = f(|X|).$$

Enumerate all triangulations reachable from X by Pachner moves while never exceeding N top-dimensional simplices. This search is finite, because there are only finitely many combinatorial triangulations with at most N top-dimensional simplices.

If T_* appears in this finite search, answer yes. If not, answer no. If $X \cong_{\text{PL}} M$, Pachner's theorem gives a Pachner path from X to T_* . The assumed height bound gives such a path staying below N , so the search finds it. If $X \not\cong_{\text{PL}} M$, no Pachner path reaches T_* , because Pachner moves preserve PL homeomorphism type.

This algorithm recognizes M , contradiction. $\square$

Corollary 27.3 (Non-recursive Pachner distance). *For the same M and T_* , there is no computable function g such that*

$$d_P(T, T_*) \leq g(|T|)$$

for every triangulation T of M .

Proof. If such a g existed, one could recognize M by enumerating all Pachner paths of length at most $g(|X|)$ from the input triangulation X , checking whether one reaches T_* . This would contradict unrecognizability of M . $\square$

Definition 27.4 (Bridge profile). *Define*

$$B_M(n) = \max\left(\{|T_*|\} \cup \{h_P(T, T_*) : T \cong_{\text{PL}} M, |T| \leq n\}\right).$$

Proposition 27.5 (Turing degree of the bridge profile). *The function B_M has the same Turing degree as the recognition problem for M .*

Proof. Given an oracle for B_M , decide whether an input triangulation X is PL homeomorphic to M as follows. Compute

$$N = B_M(|X|).$$

Enumerate all triangulations reachable from X by Pachner moves while staying below height N . If T_* appears, answer yes; otherwise answer no. If $X \cong_{\text{PL}} M$, the definition of B_M ensures that such a bounded-height path exists. If not, no path exists at all. Thus recognition of M is Turing reducible to B_M .

Conversely, assume access to an oracle recognizing M . To compute $B_M(n)$, enumerate all triangulations with at most n top-dimensional simplices. Use the oracle to select those PL homeomorphic to M . For each selected triangulation T , search over height bounds H , increasing H one by one, and perform the finite Pachner search below height H until T_* is reached. Pachner's theorem guarantees termination. This computes $h_P(T, T_*)$ for each selected T . Taking the maximum gives $B_M(n)$. $\square$

28 Application IV: Tietze Fibre Barriers

The group-theoretic analogue of the Pachner application uses the fibre of finite presentations of the trivial group.

Fix a standard decidable Tietze graph on finite group presentations, generated by elementary Nielsen–Tietze moves and generator introduction/removal moves. The moves preserve the presented group, and Tietze's theorem says that two finite presentations of isomorphic groups lie in the same connected component. Let $|P|$ denote total presentation length.

Let

$$P_* = \langle a \mid a \rangle$$

be a base presentation of the trivial group.

Definition 28.1 (Tietze height). *For a finite presentation P of the trivial group, define*

$$h_T(P, P_*) = \min_{\gamma: P \rightsquigarrow P_*} \max_{E \in \gamma} |E|,$$

where γ ranges over Tietze paths from P to P_* .

Theorem 28.2 (Non-recursive Tietze height for the trivial group). *There is no computable function*

$$f : \mathbb{N} \rightarrow \mathbb{N}$$

such that

$$h_T(P, P_*) \leq f(|P|)$$

for every finite presentation P of the trivial group.

Proof. Suppose such a computable f existed. Given an arbitrary finite presentation P , compute

$$N = f(|P|).$$

Enumerate the finite subgraph of the Tietze graph consisting of presentations of length at most N reachable from P through paths staying within that length bound. This is a finite effective search because there are only finitely many presentations of length at most N , and the chosen elementary move relation is decidable.

If P_* appears, then P presents the trivial group, since Tietze moves preserve the presented group. If P presents the trivial group, Tietze's theorem gives a path from P to P_* , and the assumed bound gives such a path staying within length N . Thus the finite search finds P_* .

This would decide whether an arbitrary finite presentation presents the trivial group, contradicting the Adian–Rabin undecidability theorem. Therefore no computable f exists. $\square$

Definition 28.3 (Tietze bridge profile). *Define*

$$B_1(n) = \max\left(\{|P_*|\} \cup \{h_T(P, P_*) : G_P \cong 1, |P| \leq n\}\right).$$

Proposition 28.4 (Turing degree of the Tietze bridge profile). *The function B_1 has the same Turing degree as the triviality problem for finitely presented groups.*

Proof. This is the move-fibre recognition theorem applied to finite group presentations, the realization map $P \mapsto G_P$, the standard Tietze graph, and the fibre over the trivial group. The hypotheses are satisfied: bounded presentations are finite and effectively enumerable, the move relation is decidable by construction, and Tietze’s theorem gives connectedness of each isomorphism fibre. $\square$

29 What Transfer Adds

The examples above share the same pattern. A theorem from one mathematical world is transported into the language of presentation access:

Trakhtenbrot undecidability	$\rightsquigarrow$	non-computable finite-model budgets,
Hilbert’s tenth problem	$\rightsquigarrow$	non-computable integer-solution heights,
undecidable theoremhood	$\rightsquigarrow$	non-computable proof-length profiles,
undecidable word problem	$\rightsquigarrow$	non-computable Dehn area profiles,
cellular-automaton nilpotency	$\rightsquigarrow$	non-computable stabilization time,
CFG ambiguity undecidability	$\rightsquigarrow$	non-computable first ambiguity length,
matrix mortality undecidability	$\rightsquigarrow$	non-computable zero-word length,
domino undecidability	$\rightsquigarrow$	non-computable finite-window obstruction radius,
$\text{MIP}^* = \text{RE}$	$\rightsquigarrow$	non-computable finite-dimensional strategy budgets,
rank-profile universality	$\rightsquigarrow$	Murphy fibres of observables,
matrix-pair similarity	$\rightsquigarrow$	wild fixed rank-profile fibres,
matroid realizability	$\rightsquigarrow$	$\exists \mathbb{R}$ -complete rank-profile realization,
fragmented positive search	$\rightsquigarrow$	non-computable observable budgets,
move-fibre recognition	$\rightsquigarrow$	non-recursive navigation in one fibre,
Adian–Rabin triviality	$\rightsquigarrow$	non-recursive Tietze height.

The transported statements expose the resource layer that is implicit in the original theorems.

This suggests a practical rule: applications should search not only for invariants, but for transfer packages. Productive hubs include incidence and matroid realization spaces, matrix-pair classification, finite-window compactness systems, computable limiting processes, nonlocal games, cellular automata, formal grammars, representation schemes, finite-poset sheaves, quiver representations with relations, matrix semigroups, Macaulay complexes, differential Galois representations, tensor networks, automata, syntactic monoids, and tropicalizations.

30 Further Directions

30.1 Transfer Atlas

The strongest current hubs are incidence-rank universality, matrix-pair classification, fragmented decidability, finite-window compactness, and computable limiting processes. Sev-

eral further hubs appear promising but require additional domain-specific input before they should be stated as theorems. Representation schemes of finitely presented groups and 3-manifold groups can import deformation-theoretic singularities into topological or group-theoretic presentation contexts. Algebraic proof systems may admit lower-bound transfer through initial ideals, tropical degenerations, or Macaulay presentations. Differential modules and Picard–Fuchs equations suggest a separate theory of scalar-compression costs for periods and parametrized integrals. Finite-volume spectral data in many-body Hamiltonian systems appears to fit the finite-window or convergence-modulus template once the effective model and promise structure are fixed carefully.

In each case, the useful question is not only whether a construction exists, but which access layer it controls: object descriptions, bounded observables, verification data, or fibre navigation.

30.2 Quantitative Yoneda Problems

Finite-test profiles such as

$$T \mapsto \text{Hom}(X, T) \quad \text{or} \quad T \mapsto \text{Hom}(T, X)$$

give budgeted analogues of Yoneda-type reconstruction. The central questions are:

- (i) how much observable budget is needed to reconstruct X ?
- (ii) how large are fibres of bounded finite-test profiles?
- (iii) when is the completeness modulus computable?
- (iv) when is it polynomial, exponential, or non-recursive?

30.3 Algebraic Verification Transfer

In proof theory, algebraic geometry, and symbolic computation, transfer packages often include maps on verification data. Potential examples include Nullstellensatz identities, Macaulay degree bounds, toric verification data, Groebner degenerations, and positivity proofs.

31 Conclusion

Presentation Theory begins with descriptions, realization maps, costs, observables, and fibres. This sequel develops the transfer calculus. A map between mathematical worlds becomes powerful when it controls access: descriptions, observable budgets, operations, verification data, and fibre geometry.

The controlled transfer schema turns this into a reusable mechanism. Every theorem written in an interpretable budgeted fragment transfers with the corresponding overheads. Fragmented positive predicates convert undecidability into non-computable resource profiles. Finite-window systems convert undecidable global existence into non-computable obstruction radii. Computable limiting systems convert undecidable threshold problems into failures of computable convergence moduli. Move-connected fibres convert recognition problems into bridge-height profiles. Rank-profile compilers convert incidence universality into Murphy-type fibres of observables. Wild-fibre compilers convert matrix-pair classification into fixed-fibre classification lower bounds.

The applications show that these mechanisms convert universality and undecidability phenomena into quantitative statements about mathematical access: singular observable fibres, wild fixed-profile fibres, $\exists\mathbb{R}$ -complete realization problems, non-computable finite-model

and Diophantine search profiles, proof-length and Dehn-area explosion, non-computable time and ambiguity profiles, non-computable finite-window and convergence-modulus bounds, non-computable observable budgets, and non-recursive navigation heights inside single fibres.

The geometry of access is itself a source of theorems.

References

- [1] S. I. Adian. The unsolvability of certain algorithmic problems in group theory. *Trudy Moskov. Mat. Obshch.* 6 (1957), 231–298.
- [2] G. R. Belitskii and V. V. Sergeichuk. Complexity of matrix problems. *Linear Algebra and its Applications* 361 (2003), 203–222.
- [3] R. Berger. The undecidability of the domino problem. *Memoirs of the American Mathematical Society* 66 (1966).
- [4] M. R. Bridson and H. Wilton. The triviality problem for profinite completions. *Inventiones Mathematicae* 202 (2015), 839–874.
- [5] M. R. Bridson. The geometry of the word problem. In *Invitations to Geometry and Topology*, Oxford Graduate Texts in Mathematics 7, Oxford University Press, 2002, 29–91.
- [6] J. Cassaigne, V. Halava, T. Harju, and F. Nicolas. Tighter undecidability bounds for matrix mortality, zero-in-the-corner problems, and more. arXiv:1404.0644.
- [7] M. Davis, H. Putnam, and J. Robinson. The decision problem for exponential Diophantine equations. *Annals of Mathematics* 74 (1961), 425–436.
- [8] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison–Wesley, 1979.
- [9] Z. Ji, A. Natarajan, T. Vidick, J. Wright, and H. Yuen. $MIP^* = RE$. arXiv:2001.04383.
- [10] J. Kari. The nilpotency problem of one-dimensional cellular automata. *SIAM Journal on Computing* 21 (1992), 571–586.
- [11] S. H. Lee and R. Vakil. Mnëv–Sturmfels universality for schemes. In *A Celebration of Algebraic Geometry*, Clay Mathematics Proceedings 18, 2013.
- [12] A. A. Markov. The insolubility of the problem of homeomorphy. In *Proceedings of the International Congress of Mathematicians*, 1958.
- [13] Yu. V. Matiyasevich. Enumerable sets are Diophantine. *Soviet Mathematics Doklady* 11 (1970), 354–358.
- [14] N. E. Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. In *Topology and Geometry–Rohlin Seminar*, Lecture Notes in Mathematics 1346, Springer, 1988, 527–543.
- [15] U. Pachner. P.L. homeomorphic manifolds are equivalent by elementary shellings. *European Journal of Combinatorics* 12 (1991), 129–145.
- [16] M. O. Rabin. Recursive unsolvability of group theoretic problems. *Annals of Mathematics* 67 (1958), 172–194.

- [17] H. Tietze. Ueber die topologischen Invarianten mehrdimensionaler Mannigfaltigkeiten. *Monatshefte fuer Mathematik und Physik* 19 (1908), 1–118.
- [18] B. A. Trakhtenbrot. The impossibility of an algorithm for the decision problem on finite classes. *AMS Translations, Series 2* 23 (1963), 1–5.
- [19] R. Vakil. Murphy’s Law in algebraic geometry: badly-behaved deformation spaces. *Inventiones Mathematicae* 164 (2006), 569–590.