

Presentation Theory V: Fibre Geometry, Navigation, and Residual Moduli

Luca Blanchi

June 2026

Abstract

This paper develops a fibre-geometric layer of Presentation Theory. A presentation system does not only assign to an object the cost of its cheapest description. It also induces a costed geometry on the kernel pair of the realization map, hence on every fibre of equivalent descriptions. The basic objects are realization fibres, cost filtrations, vertical move systems, move groupoids, height distances, navigation profiles, vertical atlases, normal-form contractions, residual moduli, entropy profiles, stabilization maps, and higher coherence data.

The formal results identify several reusable mechanisms. Height navigation is the merge level in the filtered fibre, hence satisfies an ultrametric inequality. In effective complete move systems, computable control of global fibre navigation is equivalent to computable recognition of the kernel pair; equivalently, the bridge-width profile of the finite-stage move filtration has the same Turing degree as equivalence of descriptions. Consequently, undecidable equivalence rules out computable vertical atlases and computably bounded proof-carrying normal forms. Entropy and rate-distortion profiles give lower bounds for any finite access mechanism that resolves a fibre at a declared scale. Base-residual decompositions split navigation into base motion, lifting, and residual correction. Obstruction and higher-navigation profiles measure the cost of gluing local normal forms and comparing verification paths. Stabilization may collapse fibre geometry, and vertical quasi-embeddings transfer navigation and entropy lower bounds.

The applications are included as tests of the formalism. A quantitative Novikov-Pachner theorem shows that, in dimensions at least five, there is no computable bound on the size of triangulations needed along Pachner paths from arbitrary triangulated spheres to a fixed standard sphere. A thickening bottleneck separates stable Andrews-Curtis difficulty into the cost of entering the thickenable sector and the cost of contracting inside it. A rank-fibre theorem for multiparameter persistence embeds simultaneous similarity of matrix pairs inside one rank-invariant fibre, giving residual entropy at least $r^2 \log_2 q$ over $\mathbb{F}_q$, that is, at least r -squared log-base-two q bits over the finite field with q elements. Further sections formulate CFI residual entropy for Weisfeiler-Leman type observables, coherence gaps in rewriting systems, and Markov-fibre entropy and conductance bottlenecks in algebraic statistics.

1 Introduction

A presentation system starts with descriptions, a realization map, and a cost:

$$\Gamma = (D, \rho, \kappa), \quad \rho : D \rightarrow X.$$

The usual presentation complexity of $x \in X$ is

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

This number records the cheapest description of x . It does not record the shape of the whole fibre

$$F_x = \rho^{-1}(x).$$

That fibre may be finite or infinite, connected or disconnected, easy or hard to navigate, rigid or entropic, contractible or stacky. It may contain hidden moduli, holonomy, stabilizers, nontrivial loops of verification paths, or enormous barriers between low-cost descriptions.

The thesis of this paper is that every presentation system induces not only a complexity function on objects, but a costed geometry on the kernel pair

$$K_\rho = D \times_X D = \{(d, e) \in D \times D : \rho(d) = \rho(e)\}.$$

Fibre Geometry studies this costed kernel pair. It asks when two descriptions of the same object are move-connected, how high a path must climb, how much information remains in a residual fibre, whether a normal form contracts the fibre, whether local normal forms glue, how much coherence is needed to compare verification paths, and what stabilization erases.

The guiding principle is:

every “up to” carries a hidden fibre geometry.

The paper is organized around three engines. The first is navigation: height, mountain-pass barriers, atlases, and normal-form contractions. The second is information: packing, covering, entropy, and rate-distortion of residual fibres. The third is coherence: higher paths, filling profiles, obstruction classes, and stacky normal forms.

Notation guide

The main symbols are used as follows.

Presentation system. Gamma is the triple consisting of descriptions D , a realization map ρ from descriptions to objects, and a cost function κ .

Fibre. The fibre over an object x is the set of all descriptions d whose realization is x .

Kernel pair. The kernel pair is the set of pairs of descriptions with the same realization.

Move system. M is a realization-preserving move relation between descriptions.

Height distance. H_M of d and e is the least possible maximum cost along an M -path from d to e .

Navigation profile. N of n is the worst height distance between equivalent descriptions of cost at most n .

Entropy. Ent of n and R is log-base-two of the largest R -separated set of descriptions of cost at most n inside one fibre.

Residual data. In a factorization through intermediate data, base motion explains the movement seen downstairs, while residual correction is the remaining motion inside the finer fibre.

2 Presentation systems and kernel pairs

Definition 2.1 (Presentation system). *A presentation system is a triple*

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

where D is a class or set of descriptions, $\rho : D \rightarrow X$ is a realization map, and $\kappa : D \rightarrow \Lambda$ is a cost function valued in an ordered set. Unless stated otherwise, $\Lambda = \mathbb{N}$.

Definition 2.2 (Fibre and kernel pair). *The realization fibre of $x \in X$ is*

$$F_x = \rho^{-1}(x).$$

The truncated fibre of radius n is

$$F_x^{\leq n} = \{d \in F_x : \kappa(d) \leq n\}.$$

The kernel pair is

$$K_\rho = D \times_X D = \{(d, e) : \rho(d) = \rho(e)\}.$$

The truncated kernel pair is

$$K_\rho^{\leq n} = \{(d, e) \in K_\rho : \kappa(d), \kappa(e) \leq n\}.$$

Definition 2.3 (Ground states). *The presentation complexity of x is*

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

If the infimum is attained, the ground-state fibre is

$$G_x = \{d \in F_x : \kappa(d) = C_\Gamma(x)\}.$$

Remark 2.4. *Presentation complexity is the ground-state energy of the fibre. Fibre Geometry studies the whole energy landscape, not only its minimum.*

Definition 2.5 (Effective presentation system). *A presentation system is effective if descriptions are encoded by finite strings, membership in D is decidable, κ is computable, and every ball*

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

is finite and effectively enumerable.

Example 2.6. *Finite group presentations, triangulations of closed combinatorial manifolds, integer fibres $Au = b$ in algebraic statistics, chain complexes with chosen bases, quiver representations with matrices, and filtered modules are all presentation systems once a realization map and a cost are declared.*

3 Vertical move systems

Definition 3.1 (Vertical move system). *A vertical move system on (D, ρ, κ) is a symmetric relation*

$$M \subseteq D \times D$$

such that every move preserves realization:

$$dMe \implies \rho(d) = \rho(e).$$

It is sound if this implication holds, and complete if the equivalence relation generated by M is exactly K_ρ .

Definition 3.2 (Vertical presentation system). *A vertical presentation system is a quadruple*

$$V = (D, \rho, \kappa, M),$$

where (D, ρ, κ) is a presentation system and M is a sound vertical move system.

Definition 3.3 (Path length and height). *An M -path from d to e is a sequence*

$$\gamma : d = d_0 \rightarrow d_1 \rightarrow \cdots \rightarrow d_r = e$$

with $d_i M d_{i+1}$. Its length is $|\gamma| = r$, and its height is

$$\text{ht}(\gamma) = \max_i \kappa(d_i).$$

The vertical path length and height distance are

$$L_M(d, e) = \min |\gamma|, \quad H_M(d, e) = \min \text{ht}(\gamma),$$

where the minimum is taken over M -paths from d to e , and the value is ∞ if no such path exists.

Definition 3.4 (Fibre groupoid). *The fibre groupoid $\Pi_M(\rho)$ has objects the descriptions $d \in D$ and morphisms formal M -paths, with composition by concatenation. If relations between paths are specified, one obtains a quotient groupoid or a higher groupoid.*

Definition 3.5 (Higher vertical system). *A vertical 2-presentation system is a tuple*

$$V_2 = (D, \rho, \kappa, M, R_2),$$

where R_2 is a set of elementary relations between M -paths with the same endpoints. Higher systems have cells R_i in all dimensions $i \geq 2$.

4 Height geometry

For a fixed fibre F_x , the cost function restricts to an energy function. The truncated fibres

$$F_x^{\leq 0} \subseteq F_x^{\leq 1} \subseteq F_x^{\leq 2} \subseteq \cdots$$

form a filtration of the move graph on F_x .

Definition 4.1 (Merge height). *For $d, e \in F_x$, the merge height is*

$$m_x(d, e) = \min\{h : d, e \text{ lie in the same connected component of } F_x^{\leq h}\}.$$

If d and e are not M -connected, set $m_x(d, e) = \infty$.

Theorem 4.2 (Merge-height theorem). *For every $d, e \in F_x$,*

$$H_M(d, e) = m_x(d, e).$$

Proof. If $H_M(d, e) \leq h$, there is an M -path from d to e all of whose vertices have cost at most h . Hence d and e lie in the same component of $F_x^{\leq h}$, so $m_x(d, e) \leq h$. Taking the least such h gives $m_x(d, e) \leq H_M(d, e)$.

Conversely, if $m_x(d, e) \leq h$, then d and e lie in the same connected component of the finite or infinite graph $F_x^{\leq h}$. By definition of connected component, there is an M -path between them entirely inside $F_x^{\leq h}$. Hence $H_M(d, e) \leq h$. Taking the least such h gives $H_M(d, e) \leq m_x(d, e)$. $\square$

Corollary 4.3 (Ultrametric inequality). *For d, e, f in the same M -connected fibre,*

$$H_M(d, f) \leq \max\{H_M(d, e), H_M(e, f)\}.$$

Proof. Let $h = \max\{H_M(d, e), H_M(e, f)\}$. Then d, e, f all lie in the same component of $F_x^{\leq h}$. Therefore $H_M(d, f) \leq h$. $\square$

Definition 4.4 (Mountain-pass profiles). *The mountain-pass excess between d and e is*

$$MP(d, e) = H_M(d, e) - \max\{\kappa(d), \kappa(e)\}.$$

The mountain-pass profile of a fibre is

$$MP_x(n) = \max_{d, e \in F_x^{\leq n}} MP(d, e).$$

The ground-state barrier is

$$H_{\min}(x) = \max_{d, e \in G_x} H_M(d, e),$$

when the maximum is defined.

Remark 4.5. *Height geometry is non-Archimedean. It is a geometry of barriers rather than shortest paths. The merge tree of the filtration $F_x^{\leq h}$ is the vertical analogue of a dendrogram.*

5 Navigation and recognition

Definition 5.1 (Navigation profiles). *The fibre navigation profile of x is*

$$N_x(n) = \max_{d, e \in F_x^{\leq n}} H_M(d, e).$$

The global navigation profile is

$$N_{\rho, M}(n) = \max\{H_M(d, e) : \kappa(d), \kappa(e) \leq n, \rho(d) = \rho(e)\}.$$

The pointed bridge profile for a basepoint $b_x \in F_x$ is

$$B_{x, b_x}(n) = \max_{d \in F_x^{\leq n}} H_M(d, b_x).$$

Theorem 5.2 (Kernel-pair navigation equivalence). *Let $V = (D, \rho, \kappa, M)$ be an effective vertical presentation system with decidable, sound, complete moves. Then the global navigation profile and the kernel pair have the same Turing degree:*

$$N_{\rho, M} \equiv_T K_\rho.$$

More explicitly, an oracle for $N_{\rho, M}$ decides K_ρ , and an oracle for K_ρ computes $N_{\rho, M}$.

Proof. Assume first an oracle for $N = N_{\rho, M}$. Given d, e , let $n = \max\{\kappa(d), \kappa(e)\}$ and $H = N(n)$. Enumerate the finite graph on $D^{\leq H}$ whose edges are M -moves. If e lies in the connected component of d , output yes. Soundness proves correctness of a yes answer. If e is not found and $\rho(d) = \rho(e)$, completeness gives an M -path from d to e , and the definition of $N(n)$ gives one of height at most H , contradiction. Hence no is correct.

Conversely assume an oracle for K_ρ . To compute $N(n)$, enumerate $D^{\leq n}$ and use the oracle to list all equivalent pairs in that finite ball. For each equivalent pair d, e , search $h = n, n+1, \dots$, enumerate $D^{\leq h}$, and test whether d and e are connected in the M -graph on that ball. Completeness guarantees termination. Taking the maximum of the first successful h 's gives $N(n)$. $\square$

Corollary 5.3. *If K_ρ is undecidable, then $N_{\rho, M}$ has no computable majorant.*

Proof. A computable majorant f would decide K_ρ by the first half of the preceding proof, replacing $N(n)$ by $f(n)$. $\square$

Proposition 5.4 (Move filtrations and bridge width). *Assume the move relation M is decidable and that cost balls $D^{\leq h}$ are finite and computable. Define*

$$d E_h^M e$$

to mean that d and e are connected in the finite M -graph induced on $D^{\leq h}$. Then (E_h^M) is a decidable increasing filtration of the kernel pair K_ρ . Its bridge-width profile is Turing-equivalent to K_ρ , and it has a computable majorant if and only if K_ρ is decidable.

Proof. Each E_h^M is decidable because the induced graph on the finite computable set $D^{\leq h}$ is finite. The relations are increasing in h . Soundness of moves gives $E_h^M \subseteq K_\rho$, and completeness gives

$$K_\rho = \bigcup_h E_h^M.$$

Therefore this is a filtered effective presentation system for the kernel pair. Applying the bridge-width theorem of Part I gives the stated Turing equivalence and the computable-majorant criterion. $\square$

Remark 5.5. *The navigation profile and the bridge-width profile are the same obstruction viewed with two different coordinates. Navigation emphasizes geometry inside fibres; bridge width emphasizes the finite-stage semidecision procedure for equivalence.*

6 Vertical atlases

Definition 6.1 (Vertical atlas). *A vertical chart is a domain $U \subseteq D$, a local normalizer*

$$\nu_U : U \rightarrow D, \quad \rho(\nu_U(d)) = \rho(d),$$

and a verification path $d \rightarrow \nu_U(d)$. A vertical atlas consists of chart domains U_i , local normalizers ν_i , transition paths between $\nu_i(d)$ and $\nu_j(d)$ on overlaps, and optionally higher coherence data among transitions.

Definition 6.2 (Atlas overhead). *A vertical atlas has overhead $A(n)$ if every pair d, e with $\kappa(d), \kappa(e) \leq n$ and $\rho(d) = \rho(e)$ can be connected using chart normalizations and transitions while remaining at cost at most $A(n)$.*

Theorem 6.3 (Atlas implies navigation). *If a vertical atlas has overhead $A(n)$, then*

$$N_{\rho, M}(n) \leq A(n)$$

up to the declared overhead conventions.

Proof. The atlas supplies a vertical path of height at most $A(n)$ for every equivalent pair of descriptions of cost at most n . Taking the maximum over such pairs gives the inequality. $\square$

Theorem 6.4 (No effective fibre atlas). *Let V be an effective vertical presentation system with decidable, sound, complete moves. If K_ρ is undecidable, then no computable vertical atlas can have computable overhead.*

Proof. Such an atlas would give a computable majorant for $N_{\rho, M}$, contradicting the preceding corollary. $\square$

Remark 6.5. *Undecidability is not only the absence of a classifier. It is the absence of computably controlled vertical coordinates with computably controlled transitions.*

7 Normal forms as contractions

Definition 7.1 (Point and proof-carrying normal forms). *A point normal form is a section $s : X \rightarrow D$ of ρ . A proof-carrying normal form is a section s together with, for every $d \in D$, a verification path*

$$\gamma_d : d \rightarrow s(\rho(d)).$$

Its contraction overhead is

$$C_s(n) = \max_{\kappa(d) \leq n} \text{ht}(\gamma_d).$$

Theorem 7.2 (Normal-form contraction principle). *Let s be a proof-carrying normal form. Then for every x ,*

$$C_s(n) \geq B_{x,s(x)}(n).$$

If s is used to solve arbitrary pairs in K_ρ by routing through $s(\rho(d))$, then the induced pair-routing overhead is at least $N_{\rho,M}(n)$.

Proof. For $d \in F_x^{\leq n}$, the normal-form data include a path from d to $s(x)$. Its height is at least $H_M(d, s(x))$. Maximizing over d gives the first inequality. If all equivalent pairs are routed through the same normal representative with overhead $A(n)$, then every equivalent pair of cost at most n is connected with height at most $A(n)$, so $N_{\rho,M}(n) \leq A(n)$. $\square$

Corollary 7.3. *If $N_{\rho,M}$ has no computable majorant, then no proof-carrying normal form can have computable global overhead.*

Definition 7.4 (Partial normal form). *A partial normal form is a factorization*

$$D \xrightarrow{\eta} E \xrightarrow{\sigma} X, \quad \rho = \sigma\eta.$$

It normalizes only part of the fibre. The fibres of η are residual fibres.

Remark 7.5. *Partial normal forms are often the correct structure. The residual fibre may contain unavoidable gauge, stabilizers, holonomy, or moduli.*

8 Entropy calculus

Definition 8.1 (Packing, covering, entropy). *The vertical packing number of $F_x^{\leq n}$ at scale R is*

$$\text{Pack}_x(n, R) = \max\{|S| : S \subseteq F_x^{\leq n}, H_M(d, e) > R \text{ for } d \neq e\}.$$

The vertical entropy is

$$\text{Ent}_x(n, R) = \log_2 \text{Pack}_x(n, R).$$

The covering number is

$$\text{Cover}_x(n, R) = \min\{|C| : C \subseteq F_x, \forall d \in F_x^{\leq n} \exists c \in C, H_M(d, c) \leq R\}.$$

Definition 8.2 (Rate-distortion profile). *The vertical rate-distortion profile is*

$$RD_x(n, R) = \min\{\log_2 |Q| : q : F_x^{\leq n} \rightarrow Q, \text{diam}_H(q^{-1}(a)) \leq R \text{ for all } a \in Q\}.$$

Proposition 8.3 (Packing-covering inequalities). *For every n and R ,*

$$\log_2 \text{Pack}_x(n, 2R) \leq RD_x(n, R) \leq \log_2 \text{Cover}_x(n, R),$$

up to the harmless scale convention in ultrametric balls.

Proof. If q has fibres of H -diameter at most R , then a $2R$ -separated set injects into Q . This gives the lower bound. For the upper bound, choose a cover by R -balls and map each point to a chosen centre. The fibre diameter is controlled by the ultrametric inequality, with the displayed scale convention. $\square$

Theorem 8.4 (Entropic lower bound). *Let $A : F_x^{\leq n} \rightarrow Y$ be an output mechanism. If A resolves the fibre at scale R , meaning*

$$H_M(d, e) > R \implies A(d) \neq A(e),$$

then

$$|Y| \geq \text{Pack}_x(n, R).$$

Therefore any binary encoding of A has worst-case length at least $\text{Ent}_x(n, R)$.

Proof. Take an R -separated set S of maximal size. The resolving condition makes A injective on S . Hence $|Y| \geq |S| = \text{Pack}_x(n, R)$. $\square$

Theorem 8.5 (Residual entropy after observation). *Let $O : F_x^{\leq n} \rightarrow Y$ be an observable with $|Y| \leq M$. Then some residual fibre*

$$F_{x,y}^{\leq n} = \{d \in F_x^{\leq n} : O(d) = y\}$$

contains an R -separated subset of size at least

$$\text{Pack}_x(n, R)/M.$$

Consequently

$$\text{ResEnt}_O(x, n, R) \geq \text{Ent}_x(n, R) - \log_2 M.$$

Proof. Partition a maximal R -separated set by the values of O . One part has size at least $\text{Pack}_x(n, R)/M$, and its elements remain R -separated. $\square$

Corollary 8.6 (Transcript lower bound). *If an audit or verification protocol has transcripts of length at most $\ell(n)$ over an alphabet of size q , and it resolves $F_x^{\leq n}$ at scale R , then*

$$\ell(n) \log_2 q \geq \text{Ent}_x(n, R).$$

Theorem 8.7 (Product law). *Let $V = V_1 \times V_2$ with cost*

$$\kappa(d_1, d_2) = \max\{\kappa_1(d_1), \kappa_2(d_2)\}$$

and coordinatewise moves. Then

$$H_V((d_1, d_2), (e_1, e_2)) = \max\{H_1(d_1, e_1), H_2(d_2, e_2)\}.$$

Consequently $N_V(n) = \max\{N_1(n), N_2(n)\}$, and

$$\text{Ent}_V(n, R) \geq \text{Ent}_1(n, R) + \text{Ent}_2(n, R).$$

Proof. Every product path projects to paths in both factors, so its height is at least the maximum of the factor heights. Conversely concatenate optimal height paths in the two coordinates. For entropy, the product of two R -separated sets is R -separated in the product. $\square$

9 Base-residual fibre sequences

Suppose a procedure factors as

$$D \xrightarrow{\eta} E \xrightarrow{\sigma} X, \quad \rho = \sigma\eta.$$

There is a formal vertical exact sequence

$$K_\eta \longrightarrow K_\rho \longrightarrow \eta^* K_\sigma.$$

The fibre of $K_\rho \rightarrow \eta^* K_\sigma$ over a base pair is the residual ambiguity left after the base motion.

Definition 9.1 (Lifting and residual correction). *A base path in E lifts with overhead L if every path in E of height h between $\eta(d)$ and $\eta(e)$ lifts to a path in D starting at d and having height at most $L(\kappa(d), h)$. Residual correction has overhead R if two descriptions in the same η -fibre, with the relevant costs controlled by n, h , can be connected inside that residual fibre with height at most $R(n, h)$.*

Theorem 9.2 (Base-residual navigation inequality). *Assume cost distortion $\lambda(\eta(d)) \leq a(\kappa(d))$, base navigation controlled by N_σ , base-path lifting overhead L , and residual correction overhead R . Then*

$$N_\rho(n) \leq \max\{L(n, N_\sigma(a(n))), R(n, N_\sigma(a(n)))\}$$

up to the declared cost conventions.

Proof. Let $d, e \in D$ have cost at most n and satisfy $\rho(d) = \rho(e)$. Navigate in E from $\eta(d)$ to $\eta(e)$ with height at most $N_\sigma(a(n))$. Lift that base path from d to a point d' with $\eta(d') = \eta(e)$. Then connect d' to e inside the residual η -fibre. The height is bounded by the displayed maximum. $\square$

Proposition 9.3 (Residual lower bound under no shortcuts). *Suppose the residual move geometry inside K_η embeds coarsely in the total fibre geometry, meaning residual R -separated sets remain R' -separated in K_ρ under the allowed total moves. Then residual packing and entropy lower-bound total packing and entropy at the corresponding scales.*

Proof. Under the no-shortcut hypothesis, any separated set in a residual fibre is still separated in the total fibre. Taking maxima gives the packing lower bound, and logarithms give the entropy lower bound. $\square$

Proposition 9.4 (Entropy in exact sequences). *If total fibre elements decompose into base data, residual data, and a finite set of lift choices, then*

$$\text{Ent}_\rho(n, R) \leq \text{Ent}_\sigma(a(n), R_b) + \sup_y \text{Ent}_{\eta,y}(b(n), R_r) + \text{LiftEnt}(n, R),$$

with scale conventions determined by the decomposition. Conversely, faithfully embedded base or residual entropy lower-bounds total entropy.

Proof. A cover of the base fibre, together with covers of residual fibres and representatives of lift choices, gives a cover of the total fibre. Taking logarithms yields the upper bound. The lower bounds follow by embedding separated sets. $\square$

10 Obstructions and stacky normal forms

Definition 10.1 (Local normal-form atlas). *A local normal-form atlas over a cover $\{U_i\}$ of X consists of local sections $s_i : U_i \rightarrow D$, transition paths*

$$g_{ij}(x) : s_i(x) \rightarrow s_j(x)$$

on overlaps, and optional higher coherence cells on multiple overlaps.

Definition 10.2 (Obstruction profile). *The level- r obstruction profile $\text{Obs}_r(n)$ is the least overhead needed to trivialize all r -dimensional gluing defects on the part of the system of cost at most n . If no such trivialization exists, it is ∞ .*

Theorem 10.3 (Obstruction lower bound). *If there exists an r -coherent normal-form atlas with overhead $A(n)$, then for every $i \leq r$,*

$$\text{Obs}_i(n) \leq A(n)$$

up to the cost conventions of the atlas.

Proof. An r -coherent atlas includes transitions, fillings of transition loops, fillings between fillings, and so on through level r . These are exactly the data whose minimal overhead is measured by the obstruction profiles. $\square$

Theorem 10.4 (Controlled vanishing principle). *Suppose local normal forms exist with cost $L(n)$, transitions with cost $T_1(n)$, and level- i coherences with cost $T_i(n)$ for $2 \leq i \leq r$. Then there is an r -coherent stacky normal form or r -coherent atlas with overhead bounded by*

$$\max\{L(n), T_1(n), \dots, T_r(n)\}$$

up to overhead distortions.

Proof. Use the local sections as charts, the transitions as one-dimensional gluing data, and the higher coherences to fill the cocycle defects. The maximum of the costs controls the assembled atlas. $\square$

Definition 10.5 (Holonomy). *In a factorization $D \xrightarrow{\eta} E \xrightarrow{\sigma} X$, a loop in the base fibre of σ based at $y \in E$ may lift to a transformation of the residual fibre $\eta^{-1}(y)$. The transformations generated this way form the vertical holonomy groupoid*

$$\text{Hol}_y(\eta).$$

Theorem 10.6 (Holonomy obstruction). *If vertical holonomy acts nontrivially on a residual fibre and cannot be trivialized within overhead $A(n)$, then no point-valued global normal form eliminating that residual fibre can have overhead $A(n)$.*

Proof. A point-valued global normal form chooses a representative in each residual fibre and verification paths to it. Transporting this choice around a base loop with nontrivial holonomy changes the representative by a residual transformation. To preserve a global point normal form, the data must identify the original and transported representatives within the stated overhead, contrary to the hypothesis. $\square$

Definition 10.7 (Stacky normal form). *A stacky normal form sends a description not to a single representative but to normal data together with a residual groupoid:*

$$d \mapsto (n_d, G_{\text{res}}(n_d)).$$

Principle 10.8. *Do not kill unavoidable fibre symmetry; expose it.*

11 Higher navigation

Definition 11.1 (Filling profiles). *For $r \geq 0$, an r -loop in a fibre is a combinatorial map $S^r \rightarrow F_x$ into the vertical cell complex. The r -th filling profile is*

$$\text{Fill}_r^x(n) = \max_{\alpha} \min \{ \text{cost of an } (r+1)\text{-filling of } \alpha \},$$

where α ranges over fillable r -loops of size or cost at most n . For $r = 0$, this is the navigation profile. For $r = 1$, it is a vertical Dehn function.

Theorem 11.2 (Higher navigation equivalence). *Assume an effective vertical higher presentation system with finite enumerable balls and complete decidable $(r+1)$ -cells for the intended r -loop filling problem E_r . Then*

$$\text{Fill}_r \equiv_T E_r.$$

Proof. A bound for Fill_r decides fillability by enumerating all fillings up to the bound. Conversely, an oracle for E_r identifies the fillable r -loops in the finite ball of cost at most n ; for each such loop, search until a filling is found and take the maximum of the minimal costs. $\square$

Theorem 11.3 (Higher contraction lower bound). *If a normal form is coherent through level r with overhead $A(n)$, then for every $i \leq r$,*

$$\text{Fill}_i^x(n) \leq A(n)$$

up to overhead.

Proof. A coherent contraction through level r fills all vertical i -spheres for $i \leq r$ by transporting them through the contraction to the normal point or residual groupoid, filling there, and transporting back. $\square$

Definition 11.4 (Directed vertical rewriting). *A directed vertical rewriting system is a directed move relation $d \rightarrow d'$ preserving realization. It is terminating if there are no infinite reductions, confluent if reductions from a common source can be joined, and coherently confluent through level r if the confluence data satisfy higher coherence conditions up to level r .*

Principle 11.5. *Rewriting is directed fibre contraction.*

Theorem 11.6 (Completion-to-filling bound). *Suppose a rewriting system is terminating and has completion profiles $\text{Comp}_i(n)$, where Comp_1 closes critical pairs, Comp_2 coherently relates critical-pair fillings, and so on. If the completion is coherent through level r , then*

$$\text{Fill}_i^x(n) \leq \text{Comp}_{i+1}(n)$$

for $0 \leq i \leq r-1$, up to overhead.

Proof. Confluence diagrams control point navigation. Coherent confluence diagrams fill loops of reductions. Higher coherences fill higher-dimensional defects. Taking worst-case costs gives the inequalities. $\square$

12 Stabilization

Definition 12.1 (Stabilization collapse). *A stabilization is a map $S : D \rightarrow D^+$ with $\rho^+(S(d)) = \rho(d)$. The stable height distance is*

$$H^+(d, e) = H_{M^+}(S(d), S(e)).$$

The stabilization collapse profile is

$$\text{Coll}_S(n, R) = \max\{H_M(d, e) : \kappa(d), \kappa(e) \leq n, H^+(d, e) \leq R\}.$$

The stabilization entropy collapse $\text{CollEnt}_S(n, R, r)$ is the logarithm of the largest subset of $F_x^{\leq n}$ whose points are pairwise R -separated in the original geometry but pairwise r -close after stabilization.

Theorem 12.2 (Stabilization hides entropy). *Let A be a stable invariant or stable normal form with capacity $\text{Cap}_A(n)$. If it resolves the stabilized geometry only up to scale r , then the unstable residual entropy is at least*

$$\text{CollEnt}_S(n, R, r) - \text{Cap}_A(n).$$

Proof. Take a set witnessing CollEnt_S . Its points become r -close after stabilization. A stable access mechanism of capacity $\text{Cap}_A(n)$ can distinguish at most $2^{\text{Cap}_A(n)}$ classes among them. The pigeonhole argument from residual entropy gives the bound. $\square$

Principle 12.3. *Stable equivalence may erase unstable fibre geometry.*

13 Transfer of fibre lower bounds

Definition 13.1 (Vertical morphism). *For vertical systems $V = (D, \rho, \kappa, M)$ and $W = (E, \sigma, \lambda, N)$, a vertical morphism is a pair of maps*

$$\Phi : D \rightarrow E, \quad \phi : X \rightarrow Y,$$

with $\sigma\Phi = \phi\rho$. It has cost distortion a if

$$\lambda(\Phi(d)) \leq a(\kappa(d)).$$

It is a vertical quasi-embedding if paths between images lift to paths in the source with controlled overhead.

Theorem 13.2 (Vertical transfer theorem). *Let $\Phi : V \rightarrow W$ be a vertical quasi-embedding with cost distortion a and path-lift overhead L . Then*

$$N_V(n) \leq L(n, N_W(a(n))).$$

If Φ preserves R -separation at scale $R \mapsto R'$, then

$$\text{Pack}_W(a(n), R') \geq \text{Pack}_V(n, R),$$

and hence

$$\text{Ent}_W(a(n), R') \geq \text{Ent}_V(n, R).$$

Proof. Navigate between $\Phi(d)$ and $\Phi(e)$ in W , then lift the path to V . This proves the navigation inequality. For packing, the image of an R -separated set is R' -separated by hypothesis. $\square$

14 Thermodynamic and probabilistic fibres

Definition 14.1 (Fibre types). *A fibre is vertically rigid if its low-cost part has low entropy and small navigation profile; vertically flexible if it has many descriptions but small barriers; and vertically glassy if it has both high entropy and high barriers.*

Definition 14.2 (Thermodynamic profiles). *The microcanonical fibre entropy is*

$$S_x(E) = \log |\{d \in F_x : \kappa(d) \leq E\}|.$$

The fibre partition function is

$$Z_x(\beta) = \sum_{d \in F_x} e^{-\beta \kappa(d)},$$

when the sum converges, and the free energy is

$$F_x(\beta) = -\frac{1}{\beta} \log Z_x(\beta).$$

Proposition 14.3 (Zero-temperature limit). *If the minimum $C_\Gamma(x)$ is attained and the ground-state degeneracy is finite and nonzero, then*

$$\lim_{\beta \rightarrow \infty} F_x(\beta) = C_\Gamma(x).$$

Proof. Write $c = C_\Gamma(x)$ and $g_x = |G_x|$. In the discrete cost case,

$$Z_x(\beta) = g_x e^{-\beta c} + O(e^{-\beta(c+1)}).$$

Thus $\log Z_x(\beta) = -\beta c + \log g_x + o(1)$, which gives the limit. $\square$

Proposition 14.4 (Conductance bottleneck). *Let a reversible local Markov chain on $F_x^{\leq n}$ have stationary measure π . Suppose*

$$F_x^{\leq n} = A \cup C \cup B,$$

every path from A to B passes through C , and $\pi(A), \pi(B) \geq \alpha$ while $\pi(C) \leq \varepsilon$. Under lazy bounded-degree normalization, the conductance satisfies

$$\Phi \leq \varepsilon/\alpha,$$

and the mixing time satisfies

$$\tau_{\text{mix}} \geq c\alpha/\varepsilon$$

for a universal normalization constant c .

Proof. The boundary flow of the cut separating A from the rest is bounded by the stationary mass of C , up to the transition normalization. Dividing by $\pi(A)$ gives the conductance bound, and the standard conductance lower bound gives the mixing-time bound. $\square$

15 Application I: quantitative Novikov-Pachner

Pachner's theorem says that closed combinatorial manifolds are PL homeomorphic if and only if they are connected by bistellar moves [1]. Sphere recognition is undecidable in dimensions at least five; this is one of the standard consequences of Novikov-type unrecognizability results, and appears in modern surveys of undecidable problems in topology [2].

Fix $d \geq 5$. Let D_d be the set of finite triangulations of closed combinatorial d -manifolds, let $\rho(T)$ be the PL homeomorphism type of T , and let M be the Pachner move system. Let T_0 be a fixed standard triangulation of S^d , and let $|T|$ be the number of top-dimensional simplices. Define

$$P_d(n) = \sup_{\substack{T \cong_{\text{PL}} S^d \\ |T| \leq n}} \inf_{\gamma: T \rightarrow T_0} \max_{U \in \gamma} |U|,$$

where γ ranges over Pachner paths.

Theorem 15.1 (Quantitative Novikov-Pachner). *For every $d \geq 5$, $P_d(n)$ has no computable majorant. Equivalently, for every computable function f , there is a triangulated PL d -sphere T of size at most n such that every Pachner path from T to T_0 passes through a triangulation U with*

$$|U| > f(n).$$

Proof. Suppose f were a computable majorant. Given a triangulated closed combinatorial d -manifold M of size at most n , enumerate all triangulations with at most $f(n)$ top-dimensional simplices and all Pachner moves among them. Compute the connected component of M in this finite graph. If T_0 appears, output yes; otherwise output no.

If T_0 appears, then M is PL homeomorphic to S^d . Conversely, if $M \cong_{\text{PL}} S^d$, Pachner's theorem gives a Pachner path to T_0 , and the assumed bound f gives one staying inside the enumerated finite graph. Thus the procedure decides d -sphere recognition, contradicting undecidability for $d \geq 5$. $\square$

Remark 15.2. *The theorem is a statement about the internal geometry of the single fibre of S^d . The object is simple, but navigation inside its presentation fibre has no computable height bound.*

16 Application II: thickening bottleneck

The stable Andrews-Curtis problem studies balanced presentations of the trivial group under stable Andrews-Curtis moves. Recent work of Lackenby gives explicit upper bounds for thickenable presentations of the trivial group [3]. The fibre-geometric point is the following decomposition.

Let D be the set of balanced presentations of the trivial group, let M_{st} be the stable Andrews-Curtis move system, and let $T \subseteq D$ be the thickenable sector. Let P_* be the standard presentation and $\kappa(P)$ the length of P . Define

$$A_{\text{st}}(n) = \sup_{\kappa(P) \leq n} H_{\text{st}}(P, P_*),$$

$$\Theta(n) = \sup_{\kappa(P) \leq n} \inf_{Q \in T} H_{\text{st}}(P, Q),$$

and

$$L_T(N) = \sup_{\substack{Q \in T \\ \kappa(Q) \leq N}} H_{\text{st}}(Q, P_*).$$

Theorem 16.1 (Thickening bottleneck). *Assume P_* is thickenable. Then*

$$\Theta(n) \leq A_{\text{st}}(n).$$

Moreover, if the cost of the thickenable endpoint after thickening is bounded by $E(n)$, then

$$A_{\text{st}}(n) \leq \max\{\Theta(n), L_T(E(n))\}.$$

Proof. Any path from P to P_* ends in the thickenable sector, so it is in particular a path from P to a thickenable presentation. This gives $\Theta(n) \leq A_{\text{st}}(n)$.

Conversely, choose a path from P to a thickenable Q of height at most $\Theta(n)$, with endpoint cost bounded by $E(n)$. Then connect Q to P_* inside the thickenable sector with height at most $L_T(E(n))$. Concatenation gives the displayed upper bound. $\square$

Remark 16.2. *The theorem is formal. The external mathematical input is a contraction bound for the thickenable sector. Heighted versions require explicit accounting for length growth along the moves.*

17 Application III: rank-fibre entropy

The rank invariant is complete for one-parameter persistence but incomplete in multiparameter persistence. The following construction gives a quantitative residual fibre.

Let P be the star-shaped poset with elements $p_1, \dots, p_5, t$ and relations $p_i \leq t$. Over a field k , set $V = k^r \oplus k^r$. For $A, B \in M_r(k)$, define a P -module $M(A, B)$ by $M_t = V$, $M_{p_i} = k^r$, and images

$$U_1 = k^r \oplus 0, \quad U_2 = 0 \oplus k^r, \quad U_3 = \text{graph}(I), \quad U_4 = \text{graph}(A), \quad U_5 = \text{graph}(B).$$

All maps $M_{p_i} \rightarrow M_t$ are injective.

Theorem 17.1 (Rank-fibre simultaneous similarity). *All modules $M(A, B)$ have the same rank invariant. After fixing U_1, U_2, U_3 , two modules $M(A, B)$ and $M(A', B')$ are isomorphic if and only if*

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

for some $P \in GL_r(k)$.

Proof. The rank invariant sees the dimensions of the vector spaces and the ranks along comparable pairs. All maps $p_i \leq t$ are injective from k^r to k^{2r} , so the rank invariant is constant.

An isomorphism preserving U_1 and U_2 has block diagonal form $(v, w) \mapsto (Pv, Qw)$. Preserving $U_3 = \text{graph}(I)$ forces $P = Q$. The diagonal action sends $\text{graph}(A)$ to $\text{graph}(PAP^{-1})$, and similarly for B . This is exactly simultaneous similarity. $\square$

Corollary 17.2 (Finite-field residual entropy). *Over $k = \mathbb{F}_q$, the rank-invariant fibre contains at least q^{r^2} distinct isomorphism classes. Hence its residual entropy is at least*

$$r^2 \log_2 q.$$

Proof. There are q^{2r^2} pairs (A, B) . Each simultaneous conjugacy orbit has size at most $|GL_r(\mathbb{F}_q)| < q^{r^2}$. Thus there are at least q^{r^2} orbits. $\square$

Corollary 17.3 (Residual moduli dimension). *Over an algebraically closed field, the family above contains a residual moduli quotient of generic dimension $r^2 + 1$.*

Proof. The space of pairs (A, B) has dimension $2r^2$. A generic pair has scalar stabilizer of dimension 1, so the generic orbit dimension is $r^2 - 1$. The quotient dimension is $2r^2 - (r^2 - 1) = r^2 + 1$. $\square$

18 Application IV: CFI residual entropy

The Cai-Furer-Immerman construction is a standard source of graphs indistinguishable by bounded-dimensional Weisfeiler-Leman refinement [6]. We isolate the fibre-geometric form.

Let B be a finite connected base graph. Let $C^1(B; \mathbb{F}_2)$ be the space of edge cochains and $B^1(B; \mathbb{F}_2)$ the subspace of coboundaries. Then

$$H^1(B; \mathbb{F}_2) = C^1(B; \mathbb{F}_2) / B^1(B; \mathbb{F}_2).$$

Assume a CFI-type construction $\alpha \mapsto G_\alpha$ with two properties:

- (i) for rigid bases, $G_\alpha \cong G_\beta$ if and only if $\alpha - \beta \in B^1(B; \mathbb{F}_2)$;
- (ii) for a fixed WL budget k and bases in the corresponding high-width regime, the k -WL profile is constant on the cohomological family.

Theorem 18.1 (CFI residual fibre entropy). *Under these hypotheses, one residual fibre of the k -WL observable contains at least*

$$|H^1(B; \mathbb{F}_2)| = 2^{\beta_1(B)}$$

pairwise nonisomorphic graphs. Thus

$$\text{ResEnt}_{k\text{-WL}}(B) \geq \beta_1(B).$$

Proof. Distinct cohomology classes give nonisomorphic graphs by the classification hypothesis. The WL-local-blindness hypothesis puts all of them in one k -WL residual fibre. Taking logarithms gives the entropy bound. $\square$

Corollary 18.2. *For families of rigid bases with $\beta_1(B) = \Omega(|V(B)|)$ satisfying the WL-local-blindness condition for fixed k , the k -WL observable leaves residual entropy $\Omega(|V(B)|)$.*

19 Application V: coherence gaps in rewriting

Higher-dimensional rewriting studies not only termination and confluence, but coherence among reductions. Squier-type theory connects finite derivation type and homological finiteness properties [7].

Theorem 19.1 (Coherence gap). *Let C be a finite 2-complex with Dehn function $f(n)$. There is a vertical 2-presentation system V_C such that every description is connected to a base description by a path of height $O(\kappa(d))$, while the vertical 2-filling profile contains a family of loops with filling cost equivalent to $f(n)$.*

Proof. Include a copy of C inside one realization fibre, based at b . For each additional description vertex d of cost n , attach an edge from d to b of height $O(n)$, and add no 2-cells that fill nontrivial loops in the copy of C . Then every vertex navigates to b with linear height. Loops contained in C retain their original filling lower bounds, while the original 2-cells of C still provide the corresponding upper bounds inside that subcomplex. $\square$

Corollary 19.2. *A normal form for V_C that is coherent at level 2 must have coherence overhead at least the displayed filling profile.*

Remark 19.3. *Normalization can be easy while proof equality is hard. Point navigation and coherent verification are different resources.*

20 Application VI: Markov fibres

Let $A : \mathbb{Z}^m \rightarrow \mathbb{Z}^r$ be an integer matrix and $D = \mathbb{N}^m$. Define $\rho(u) = Au$. The fibre over b is

$$F_b = \{u \in \mathbb{N}^m : Au = b\}.$$

A move is an integer vector $z \in \ker_{\mathbb{Z}} A$, applied as $u \mapsto u + z$ when $u + z \in \mathbb{N}^m$.

Definition 20.1 (Markov basis). *A finite set $M \subseteq \ker_{\mathbb{Z}} A$ is a Markov basis if the corresponding moves connect every fibre F_b .*

Proposition 20.2. *A finite move set M is a Markov basis if and only if it is a finite complete sound move system for the kernel pair of $\rho : \mathbb{N}^m \rightarrow \mathbb{Z}^r$.*

Proof. Soundness is $A(u + z) = Au$. Completeness is exactly the Markov-basis property that any two points in the same fibre are connected by moves from M while staying in $\mathbb{N}^m$. This is the finite-fibre form of the fundamental theorem of Markov bases [4]. $\square$

Theorem 20.3 (Markov fibre entropy lower bound). *Let $F_b^{\leq n}$ be a truncated Markov fibre with height distance induced by a Markov basis. If an invariant, compressed witness, or sampling summary resolves $F_b^{\leq n}$ at scale R and has output set Y , then*

$$\log_2 |Y| \geq \text{Ent}_b(n, R).$$

Proof. This is the entropic lower bound applied to the Markov fibre. $\square$

Theorem 20.4 (Markov sampler bottleneck). *For a reversible local Markov chain on $F_b^{\leq n}$, if*

$$F_b^{\leq n} = A \cup C \cup B$$

with every path from A to B passing through C , $\pi(A), \pi(B) \geq \alpha$, and $\pi(C) \leq \varepsilon$, then

$$\tau_{\text{mix}} \geq c\alpha/\varepsilon$$

under standard lazy-chain normalization.

Proof. Apply the conductance bottleneck proposition to the Markov fibre graph. $\square$

21 Summary

The theory above isolates three engines.

The navigation engine studies N , H , MP , vertical atlases, and normal-form contractions. Its core theorem is that effective navigation and kernel-pair recognition are Turing equivalent under decidable complete moves.

The entropy engine studies Pack, Cover, RD , Ent, residual entropy, and thermodynamic profiles. Its core theorem is that residual entropy is at least fibre entropy minus observable capacity.

The coherence engine studies Fill_r , obstruction profiles, holonomy, stacky normal forms, and coherent rewriting. Its core theorem is that coherent normal forms must pay the corresponding filling and obstruction costs.

Together they give the access inequality:

$$\text{complete access} \geq \text{navigation} + \text{information} + \text{coherence} + \text{obstruction}.$$

22 Further directions

The filtration $F_x^{\leq h}$ gives persistent fibre homology. PH_0 is the merge tree; PH_1 records persistent loops of verification paths; higher groups record higher coherence obstructions.

For towers $D \rightarrow E \rightarrow X$, one expects fibre spectral sequences of the form

$$E_{p,q}^2 = H_p^{\text{vert}}(K_\sigma; H_q^{\text{vert}}(K_\eta)) \Rightarrow H_{p+q}^{\text{vert}}(K_\rho),$$

under suitable fibration hypotheses.

Stabilization gaps measure what stable classification erases. Probabilistic fibre geometry compares worst-case and typical navigation. Fibre novelty can be measured by

$$Nov_S(d) = \min_i H(d, d_i)$$

for a known set of descriptions $S = \{d_1, \dots, d_m\}$. A later observable theory should study how observables reduce residual entropy and residual navigation complexity.

References

- [1] U. Pachner, *P.L. homeomorphic manifolds are equivalent by elementary shellings*, European Journal of Combinatorics 12, 1991.
- [2] B. Poonen, *Undecidable problems: a sampler*, Interpreting Goedel: Critical Essays, Cambridge University Press, 2014.
- [3] M. Lackenby, *The stable Andrews-Curtis conjecture and thickenable presentations of the trivial group*, arXiv:2606.06122, 2026.
- [4] P. Diaconis and B. Sturmfels, *Algebraic algorithms for sampling from conditional distributions*, Annals of Statistics 26, 1998.
- [5] G. Carlsson and A. Zomorodian, *The theory of multidimensional persistence*, Discrete and Computational Geometry 42, 2009.
- [6] J. Cai, M. Fuerer, and N. Immerman, *An optimal lower bound on the number of variables for graph identification*, Combinatorica 12, 1992.
- [7] C. C. Squier, *Word problems and a homological finiteness condition for monoids*, Journal of Pure and Applied Algebra 49, 1987.