

Presentation Theory I: Description Systems, Costs, and Observable Fibres

Luca Blanchi

June 2026

Abstract

This note develops a formal language for studying mathematical objects through chosen systems of description. A presentation system consists of descriptions, a realization map, and resource bounds; it induces bounded parts, presentation complexity, and realization fibres. In the effective setting, a filtered presentation system has a decidable description language, effectively finite cost balls, and a filtered semidecision procedure for equivalence of descriptions. For such systems the bounded-presentation problem, the minimal-cost function, and the width profile for finding low-cost representatives have the same Turing degree:

$$P_\Gamma \equiv_T C_\Gamma \equiv_T W_\Gamma.$$

Similarly, equivalence of descriptions has the same Turing degree as the bridge-width profile:

$$EQ_\Gamma \equiv_T B_\Gamma.$$

Thus computable majorants for the relevant width profiles are exactly the effective content needed to decide bounded presentation or equivalence. The theory also treats observables and costed observables, whose fibres measure indistinguishability and whose bounded profiles give lower-bound criteria. Cost-controlled maps and normal-form compilers provide transfer principles between presentation systems. A central section collects reusable abstract theorems, including bounded-image obstructions, observable factorization, normal-form transfer, verification-length barriers, and no-resolution results, before the final examples specialize the language to heights, Kolmogorov complexity, circuits, group presentations, and algebraic complexity.

1 Guiding idea

Many mathematical objects come with natural ways of being presented: by coordinates, generators and relations, equations, circuits, words, modules, kernels, local models, proof data, or finite approximations.

A presentation theory studies four related questions.

First, which descriptions are allowed?

Second, which observations are allowed?

Third, what fibres are created by realization maps and by observables?

Fourth, how do costs on descriptions or observations make these questions quantitative?

The basic principle is:

presentations realize objects, observables forget information,
and costs turn both processes into filtrations.

Conversely, an observation incompatible with every low-cost object proves that the object lies outside the corresponding bounded presentation class.

Complexity is assigned relative to a specified presentation system.

Presentation, observation, and cost play distinct roles. A description produces or specifies an object. An observable extracts information from an object. A cost imposes a budget on descriptions or observations.

This gives three complementary layers: presentation complexity for objects, observable complexity for tests, and fibre structure for the information forgotten by a chosen family of observables. In many examples, all observables of bounded arity, degree, rank, or test size factor through a smaller structural statistic.

2 Object classes

An **object class** is a class

$$\mathcal{X}$$

of mathematical objects.

Depending on the application, $\mathcal{X}$ may be a set of objects, a set of isomorphism classes, or a category/groupoid whose objects are studied up to an equivalence relation.

We write

$$\equiv_{\mathcal{X}}$$

for the chosen equivalence relation on $\mathcal{X}$. If $\mathcal{X}$ is already a set of equivalence classes, then $\equiv_{\mathcal{X}}$ is equality. In categorical or groupoid settings it may be isomorphism or another declared equivalence.

Maps between object classes are assumed to respect the declared equivalence relations.

Examples include:

$$\mathcal{X} = \mathbb{P}^n(\mathbb{Q}),$$

$$\mathcal{X} = \text{finite groups up to isomorphism,}$$

$$\mathcal{X} = \text{Boolean functions on } n \text{ variables,}$$

$$\mathcal{X} = \text{polynomials over a fixed field.}$$

The first step is to choose what counts as the object class and what equivalence relation is being imposed.

This choice is not cosmetic. A finite group considered as an abstract group, a marked group, a permutation group, a matrix group, or a group equipped with structural data may lead to different presentation systems and different costs.

3 Resource scales

A **resource scale** is a preordered set

$$(B, \leq)$$

whose elements represent possible resource bounds.

The simplest case is

$$B = \mathbb{N}$$

with the usual order.

More generally, B may be a multi-resource scale such as

$$B = \mathbb{N}^k$$

with coordinatewise order.

This allows one to record several resources at once, such as size, depth, height, degree, arity, number of generators, number of relations, or total relation length.

A **cost bound** is an element

$$b \in B.$$

Informally, b is the available budget.

Costs may be scalar or multi-resource. In many applications the bounded part is defined by several simultaneous constraints.

4 Presentation systems

A **presentation system** for an object class $\mathcal{X}$ consists of 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.

An element

$$d \in \mathcal{D}$$

is a description. The object realized by d is

$$\rho(d).$$

The cost of the description is

$$\kappa(d).$$

The presentation system tells us which descriptions are allowed and how expensive they are.

The same object class may admit many different presentation systems. This plurality is central to the theory.

For example, finite groups may be presented by generators and relations, multiplication tables, permutation actions, matrix representations, semidirect-product data, module data, commutator tensors, normal covers, or finite local models.

Different presentation systems measure different mathematical resources.

Throughout, when an object class is studied up to $\equiv_{\mathcal{X}}$, the notation

$$\rho(d) = x$$

means $\rho(d) \equiv_{\mathcal{X}} x$. This convention keeps the notation light while allowing the same formalism to cover equality of functions, isomorphism of groups, and equivalence of structured objects.

5 Admissibility and naturality

The formal definition of a presentation system is intentionally broad. In applications it is paired with admissibility conditions appropriate to the mathematical problem.

Common conditions include:

- (i) the realization map is effectively computable, verifiable, or mathematically canonical;
- (ii) bounded description classes are finite, compact, constructible, recursively enumerable, or otherwise controllable;
- (iii) the cost is monotone under natural enlargements of descriptions;
- (iv) equivalent descriptions of the same object have controlled effects on cost;
- (v) natural operations in the domain admit cost-controlled transfers when such transfers are used;
- (vi) low-cost descriptions do not encode target objects by ad hoc labels unless such labels are part of the declared presentation system;
- (vii) observables used for lower bounds are canonical, or their selection and evaluation costs are explicitly charged.

These conditions keep the formalism aligned with the intended mathematical presentation problem. They also make clear that the generality of the definition is not a license to choose arbitrary encodings after the target object is known.

The corresponding methodological convention is:

costs are attached to declared description or observation systems,
and lower-bound observables are either canonical or explicitly costed.

6 Effectively finitary presentation systems

In effective applications one often works with presentation systems whose bounded parts can actually be searched.

A presentation system

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

is **effectively finitary** if, for every cost bound b ,

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

is finite and effectively enumerable, and the relevant realization comparisons

$$\rho(d) \equiv_{\mathcal{X}} x$$

are decidable for the bounded descriptions under consideration.

This condition is deliberately modest. It says that bounded search is meaningful once the bound is fixed. It does not say that the unbounded realization problem, equivalence problem, or fibre recognition problem is decidable.

Proposition 6.1 (Bounded membership in effectively finitary systems). *Let Γ be effectively finitary. In any representation in which the comparisons $\rho(d) \equiv_{\mathcal{X}} x$ are decidable for $d \in \mathcal{D}^{\leq b}$, the predicate*

$$x \in \mathcal{X}_{\Gamma}^{\leq b}$$

is decidable.

Proof. Enumerate the finite set $\mathcal{D}^{\leq b}$ and test whether any enumerated description realizes an object equivalent to x . The search terminates because the bounded description set is finite. $\square$

7 Filtered effective presentation systems

The previous notion of effective finitariness says that bounded search is meaningful once a bound is fixed. A sharper computability model records how equality of represented objects is witnessed by finite stages.

Fix a finite alphabet Σ . A **filtered effective presentation system** is a triple

$$\Gamma = (D, \kappa, (E_t)_{t \in \mathbb{N}})$$

such that:

- (i) $D \subseteq \Sigma^*$ is decidable;
- (ii) $\kappa : D \rightarrow \mathbb{N}$ is computable and effectively proper, meaning that

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

is finite and computable uniformly in b ;

- (iii) each relation $E_t \subseteq D \times D$ is decidable uniformly in t ;
- (iv) $E_t \subseteq E_{t+1}$ for all t ;
- (v) the union

$$E = \bigcup_{t \in \mathbb{N}} E_t$$

is an equivalence relation on D .

The represented object class is the quotient D/E . The realization map is the quotient map $d \mapsto [d]$.

The **bounded-presentation problem** is

$$P_{\Gamma} = \{(d, b) \in D \times \mathbb{N} : \exists e \in D_{\leq b} \text{ with } d E e\}.$$

The minimal presentation cost of the object represented by d is

$$C_{\Gamma}(d) = \min\{\kappa(e) : e E d\}.$$

For a positive instance $(d, b) \in P_\Gamma$, define its presentation width by

$$w_\Gamma(d, b) = \min\{t : \exists e \in D_{\leq b} \text{ with } d E_t e\}.$$

The global bounded-presentation width profile is

$$W_\Gamma(N) = \max\{w_\Gamma(d, b) : |d| + b \leq N, (d, b) \in P_\Gamma\},$$

with value 0 if the set is empty.

Similarly, define the equivalence problem

$$EQ_\Gamma = \{(d, e) \in D \times D : d E e\}.$$

For $d E e$, define

$$\beta_\Gamma(d, e) = \min\{t : d E_t e\},$$

and define the bridge-width profile

$$B_\Gamma(N) = \max\{\beta_\Gamma(d, e) : |d| + |e| \leq N, d E e\},$$

again with value 0 if the set is empty.

Theorem 7.1 (Degree conservation for bounded presentation). *For every filtered effective presentation system,*

$$P_\Gamma \equiv_T C_\Gamma \equiv_T W_\Gamma.$$

Proof. The equivalence $P_\Gamma \equiv_T C_\Gamma$ is immediate from the definition. Given C_Γ , decide $(d, b) \in P_\Gamma$ by testing $C_\Gamma(d) \leq b$. Given an oracle for P_Γ , compute $C_\Gamma(d)$ by testing $b = 0, 1, \dots, \kappa(d)$ and taking the least b with $(d, b) \in P_\Gamma$.

Now suppose W_Γ is available. Given (d, b) , set $N = |d| + b$ and $M = W_\Gamma(N)$. Since $D_{\leq b}$ is finite and computable, one can decide whether

$$\exists e \in D_{\leq b} \quad d E_M e.$$

If such e exists, then $(d, b) \in P_\Gamma$. If no such e exists and (d, b) were positive, then by definition $w_\Gamma(d, b) \leq W_\Gamma(N) = M$, contradiction. Hence $P_\Gamma \leq_T W_\Gamma$.

Conversely, assume an oracle for P_Γ . To compute $W_\Gamma(N)$, list the finitely many pairs (d, b) with $d \in D$ and $|d| + b \leq N$. Use the oracle to keep exactly the positive pairs. For each positive pair, search through $t = 0, 1, 2, \dots$ until some $e \in D_{\leq b}$ satisfies $d E_t e$. This search terminates by positivity. Taking the maximum of the finitely many widths gives $W_\Gamma(N)$. Thus $W_\Gamma \leq_T P_\Gamma$. $\square$

Theorem 7.2 (Degree conservation for equivalence). *For every filtered effective presentation system,*

$$EQ_\Gamma \equiv_T B_\Gamma.$$

Proof. Given B_Γ , decide whether $d E e$ by setting $N = |d| + |e|$, computing $M = B_\Gamma(N)$, and testing the decidable relation $d E_M e$. If $d E e$, then $\beta_\Gamma(d, e) \leq M$, so the test succeeds; if the test succeeds, then $d E e$ because $E_M \subseteq E$.

Conversely, given an oracle for EQ_Γ , list all pairs with $|d| + |e| \leq N$, keep the equivalent pairs, compute each $\beta_\Gamma(d, e)$ by finite-stage search, and take the maximum. This computes $B_\Gamma(N)$. $\square$

Corollary 7.3 (Computable majorants). *The following are equivalent:*

$$P_\Gamma \text{ is decidable,} \quad C_\Gamma \text{ is computable,} \quad W_\Gamma \text{ has a computable majorant.}$$

Likewise,

$$EQ_\Gamma \text{ is decidable} \iff B_\Gamma \text{ has a computable majorant.}$$

Proof. If the decision problem is decidable, the corresponding width profile is computable by the preceding proofs and hence majorizes itself. Conversely, suppose g is a computable majorant for W_Γ . To decide $(d, b) \in P_\Gamma$, set $N = |d| + b$ and test whether some $e \in D_{\leq b}$ satisfies $d E_{g(N)} e$. The same argument used in the proof of degree conservation shows that this test is correct. The bridge-width statement is identical, using $g(|d| + |e|)$ in place of $B_\Gamma(|d| + |e|)$. $\square$

Proposition 7.4 (Cofinal filtrations). *Let (E_t) and (F_s) be two decidable increasing filtrations of the same equivalence relation E on D . Suppose they are computably cofinal: there are computable functions $\alpha, \beta : \mathbb{N} \rightarrow \mathbb{N}$ such that*

$$E_t \subseteq F_{\alpha(t)}, \quad F_s \subseteq E_{\beta(s)}.$$

Then the corresponding W -profiles and B -profiles majorize each other up to computable change of variable. In particular, the existence of a computable majorant is independent of replacing a filtration by a computably cofinal one.

Proof. If $E_t \subseteq F_{\alpha(t)}$, let

$$\alpha^*(m) = \max_{0 \leq t \leq m} \alpha(t).$$

Then every positive bounded-presentation instance witnessed by the E -filtration by stage m is witnessed by the F -filtration by stage $\alpha^*(m)$. Hence $W_F(N) \leq \alpha^*(W_E(N))$. The other inequality follows from β . The proof for bridge width is the same. $\square$

Proposition 7.5 (Sharpness of the two degrees). *Every c.e. Turing degree occurs as $\deg_T(P_\Gamma)$ for some filtered effective presentation system. More generally, for every pair of c.e. degrees $\mathbf{a} \leq \mathbf{b}$, there is a filtered effective presentation system with*

$$\deg_T(P_\Gamma) = \mathbf{a}, \quad \deg_T(EQ_\Gamma) = \mathbf{b}.$$

Proof. For a c.e. set $A = \bigcup_t A_t$, with A_t finite and decidable uniformly in t , use descriptions $0x$ and $1x$. Set

$$\kappa(0x) = |x|, \quad \kappa(1x) = |x| + 1,$$

and let $0x$ and $1x$ become equivalent at stage t precisely when $x \in A_t$. Then

$$x \in A \iff (1x, |x|) \in P_\Gamma,$$

and the equivalence problem has the same degree.

For a pair $\mathbf{a} \leq \mathbf{b}$, choose c.e. sets A, B with degrees $\mathbf{a}, \mathbf{b}$ and $A \leq_T B$. Take the disjoint union of the preceding A -system with a second component whose nontrivial equivalences encode B but whose equivalent pairs have equal cost. The second component contributes B to equivalence but only a decidable bounded-presentation problem. The bounded-presentation degree is therefore $\mathbf{a}$, while the equivalence degree is $\mathbf{b}$. $\square$

8 Bounded descriptions and bounded objects

For a cost bound $b \in B$, define the bounded description class

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

The corresponding bounded part of $\mathcal{X}$ is

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

Thus

$$\mathcal{X}_{\Gamma}^{\leq b}$$

is the class of objects admitting at least one Γ -description of cost at most b .
Equivalently, an object $x \in \mathcal{X}$ is b -presentable if

$$x \in \mathcal{X}_{\Gamma}^{\leq b}.$$

If $b \leq b'$, then

$$\mathcal{X}_{\Gamma}^{\leq b} \subseteq \mathcal{X}_{\Gamma}^{\leq b'}.$$

Hence the bounded parts form a filtration of $\mathcal{X}$:

$$\mathcal{X}_{\Gamma}^{\leq b} \subseteq \mathcal{X}_{\Gamma}^{\leq b'} \quad \text{whenever } b \leq b'.$$

This is the **bounded-presentation filtration** associated with Γ .

The filtration is often the primary object. A single numerical complexity value may be less informative than the entire family of bounded parts.

9 Presentation complexity

When the resource scale admits a meaningful notion of infimum or minimum, one may define the **presentation complexity** of an object $x \in \mathcal{X}$ by

$$C_{\Gamma}(x) = \inf\{\kappa(d) : d \in \mathcal{D}, \rho(d) \equiv_{\mathcal{X}} x\}.$$

In discrete settings this is often a minimum:

$$C_{\Gamma}(x) = \min\{\kappa(d) : \rho(d) \equiv_{\mathcal{X}} x\}.$$

Thus $C_{\Gamma}(x)$ is the cost of the cheapest allowed description of x .

The bounded part can then be written as

$$\mathcal{X}_{\Gamma}^{\leq b} = \{x \in \mathcal{X} : C_{\Gamma}(x) \leq b\},$$

whenever the minimum is attained or the inequality is interpreted as the existence of a description of cost at most b .

The primary object is often the full filtration

$$b \mapsto \mathcal{X}_{\Gamma}^{\leq b},$$

with $C_{\Gamma}(x)$ as its pointwise numerical summary when such a summary is available.

10 Basic upper bound principle

If $d \in \mathcal{D}$ is a description of x , meaning

$$\rho(d) = x,$$

then

$$C_{\Gamma}(x) \leq \kappa(d).$$

Thus every explicit description gives an upper bound on presentation complexity.
This is the constructive half of the theory.

11 Presentation fibres

The realization map

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

is not only a way to assign objects to descriptions. It is a map with fibres.

For $x \in \mathcal{X}$, define the **presentation fibre**

$$\mathfrak{F}_\rho(x) = \rho^{-1}(x) = \{d \in \mathcal{D} : \rho(d) \equiv_{\mathcal{X}} x\}.$$

Thus $\mathfrak{F}_\rho(x)$ is the class of all descriptions realizing x .

The numerical invariant $C_\Gamma(x)$ sees only the cheapest point in this fibre:

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

The full fibre may contain many other descriptions, with their own logic, computability, and geometry. A small value of $C_\Gamma(x)$ says that the fibre meets a low-cost region of $\mathcal{D}$; it does not by itself determine how difficult the fibre is to recognize, normalize, or navigate.

For a cost bound b , the truncated fibre is

$$\mathfrak{F}_\rho^{\leq b}(x) = \{d \in \mathfrak{F}_\rho(x) : \kappa(d) \leq b\}.$$

Studying the family

$$b \mapsto \mathfrak{F}_\rho^{\leq b}(x)$$

refines the study of $C_\Gamma(x)$. It records how the descriptions of a fixed object populate the bounded parts of the description space.

12 Fibre membership and verification

In effective settings, assume the description class is encoded by finite words and let

$$\ell : \mathcal{D} \rightarrow \mathbb{N}$$

be a coding length. For a fixed object x , the **fibre recognition problem** is the decision problem

$$d \longmapsto [\rho(d) = x].$$

This problem may be much harder than producing one economical description of x .

A **verification relation** for the fibre of x is a decidable relation

$$R_x(d, \pi)$$

such that

$$\rho(d) = x \iff \exists \pi R_x(d, \pi).$$

The datum π is membership verification data for the statement $d \in \mathfrak{F}_\rho(x)$.

Define the minimum verification-data length

$$c_{R_x}(d) = \min\{|\pi| : R_x(d, \pi)\},$$

for $d \in \mathfrak{F}_\rho(x)$, and the bounded verification-length profile

$$L_{R_x}(n) = \max_{\substack{\ell(d) \leq n \\ d \in \mathfrak{F}_\rho(x)}} c_{R_x}(d),$$

with the usual convention that the maximum over an empty set is 0.

This profile measures the length of the shortest membership verification data needed to verify all fibre elements of coding length at most n .

13 Move systems on fibres

A **move system** on descriptions is a relation

$$d \sim_{\mathcal{M}} e$$

on $\mathcal{D}$, usually chosen so that moves preserve realization:

$$d \sim_{\mathcal{M}} e \implies \rho(d) = \rho(e).$$

Then every presentation fibre becomes a graph. Its vertices are the descriptions in $\mathfrak{F}_\rho(x)$, and its edges are the elementary moves. In this note, bridge profiles use symmetric move systems; directed move systems can be treated by choosing an orientation convention for paths.

Let

$$d_{\mathcal{M}}(d, e)$$

be the shortest path length from d to e , with value ∞ if no path exists.

After fixing a base description $d_x \in \mathfrak{F}_\rho(x)$, define the **bridge profile**

$$B_{\rho, \mathcal{M}, x, d_x}(n) = \sup_{\substack{\ell(d) \leq n \\ d \in \mathfrak{F}_\rho(x)}} d_{\mathcal{M}}(d, d_x).$$

This measures how hard it is to return bounded descriptions of x to a chosen base description using the allowed moves.

A complementary invariant is the packing profile

$$\text{Pack}_{\rho, \mathcal{M}}(x; n, R) = \sup\{|S| : S \subseteq \mathfrak{F}_\rho(x), \ell(d) \leq n \ \forall d \in S, d_{\mathcal{M}}(d, e) > R \ \forall d \neq e\}.$$

The bridge profile measures long normalization paths; the packing profile measures many mutually distant descriptions. Diameter and entropy-type quantities can be recovered from these by taking pairwise suprema or logarithms of packing numbers.

14 Running example: group presentations

Finite presentations of groups give a useful model for the preceding definitions. Descriptions are finite presentations

$$\langle S \mid R \rangle,$$

and the realization map sends such a presentation to the group

$$F(S)/\langle\langle R \rangle\rangle,$$

viewed up to isomorphism. A natural cost records quantities such as

$$(|S|, |R|, \sum_{r \in R} |r|).$$

The presentation fibre of a group G is the class of all finite presentations realizing a group isomorphic to G . Tietze transformations form a natural reversible move system on this fibre, and the associated bridge profile measures how hard it is to connect a bounded presentation of G to a chosen base presentation.

This example also shows why presentation and observation must be kept separate. Words used as relators are part of a description and contribute to presentation cost. Words used as word maps, word equations, or counting tests are observables and contribute to observation cost.

15 Cost-controlled maps

Let

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

be a presentation system for $\mathcal{X}$, and let

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

be a presentation system for $\mathcal{Y}$.

Let

$$F : \mathcal{X} \rightarrow \mathcal{Y}$$

be a map of object classes.

A **description-level transfer** for F is a map

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

such that

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

for every $d \in \mathcal{D}_{\mathcal{X}}$.

The transfer is **cost-controlled with monotone overhead** $f : B_{\mathcal{X}} \rightarrow B_{\mathcal{Y}}$ if

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

for every $d \in \mathcal{D}_{\mathcal{X}}$.

The overhead function f is part of the declared transfer data. It need not be optimal, but it must be monotone and compatible with the chosen source and target resource scales. In effective settings one usually requires f to be computable; in quantitative complexity settings one may restrict to linear, polynomial, or primitive recursive overheads.

In that case,

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

When presentation complexities are defined and the minima or infima are compatible with the overhead, this gives

$$C_{\Gamma_Y}(F(x)) \leq f(C_{\Gamma_X}(x)).$$

Thus cost-controlled transfers formalize the idea that natural operations transport low-cost descriptions to low-cost descriptions.

16 Cost-controlled operations

More generally, let

$$F : \mathcal{X}_1 \times \cdots \times \mathcal{X}_n \rightarrow \mathcal{Y}$$

be an (n) -ary operation.

A description-level transfer for F is a map

$$\hat{F} : \mathcal{D}_1 \times \cdots \times \mathcal{D}_n \rightarrow \mathcal{D}_Y$$

such that

$$\rho_Y(\hat{F}(d_1, \dots, d_n)) = F(\rho_1(d_1), \dots, \rho_n(d_n)).$$

It is cost-controlled with monotone overhead f if

$$\kappa_Y(\hat{F}(d_1, \dots, d_n)) \leq f(\kappa_1(d_1), \dots, \kappa_n(d_n)).$$

Then

$$C_{\Gamma_Y}(F(x_1, \dots, x_n)) \leq f(C_{\Gamma_1}(x_1), \dots, C_{\Gamma_n}(x_n))$$

whenever the corresponding presentation complexities are defined.

This is the basic form of a **cost calculus**.

Examples include products, quotients, extensions, tensor products, composition of maps, base change, blow-ups, reductions modulo primes, and passage to associated graded or local models.

17 Simulation between presentation systems

Let

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

and

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

be presentation systems for the same object class $\mathcal{X}$.

A **translation** from Γ to Δ with monotone overhead f is a map

$$T : \mathcal{D}_\Gamma \rightarrow \mathcal{D}_\Delta$$

such that

$$\rho_\Delta(T(d)) = \rho_\Gamma(d)$$

and

$$\kappa_\Delta(T(d)) \leq f(\kappa_\Gamma(d)).$$

Then

$$\mathcal{X}_\Gamma^{\leq b} \subseteq \mathcal{X}_\Delta^{\leq f(b)}.$$

When numerical presentation complexity is defined,

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

Thus Δ simulates Γ up to monotone overhead f .

If there are translations in both directions with controlled overheads, then the two presentation systems are equivalent up to the corresponding resource distortion.

Separations between presentation systems arise when a family (x_n) has low complexity in one system and high complexity in another.

18 Counting and volume

The bounded parts define a counting function whenever cardinality is meaningful:

$$b \mapsto |\mathcal{X}_\Gamma^{\leq b}|.$$

This is the **volume profile** of the presentation system.

Counting arguments give generic lower bounds. If a family $\mathcal{S} \subseteq \mathcal{X}$ satisfies

$$|\mathcal{S}| > |\mathcal{D}^{\leq b}|,$$

then some $x \in \mathcal{S}$ has no description of cost at most b .

This follows from

$$|\mathcal{X}_\Gamma^{\leq b}| \leq |\mathcal{D}^{\leq b}|.$$

Similarly, if

$$|\omega(\mathcal{S})| > |\text{Val}_{\Gamma, \omega}(b)|,$$

then some $x \in \mathcal{S}$ has no description of cost at most b .

Thus volume and observable volume provide nonconstructive lower bounds.

Counting arguments prove the existence of complex objects. Additional structure turns those existence statements into explicit examples.

19 Observables

Let $\mathcal{X}$ be an object class.

An **observable** on $\mathcal{X}$ is a map

$$\omega : \mathcal{X} \rightarrow A$$

to an observation space A .

The value

$$\omega(x)$$

is the information about x seen by the observable.

The observation space A need not be numerical. It may consist of numbers, vectors, modules, lattices, distributions, ranks, graphs, isomorphism classes, hom-counts, reductions modulo primes, local data, cohomology groups, or other structured data.

A cost measures the expense of a description. An observable measures information extracted from the object.

A numerical invariant may serve as an observable independently of any presentation-cost interpretation.

20 Attainable observations below a bound

Let

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

be a presentation system for $\mathcal{X}$, and let

$$\omega : \mathcal{X} \rightarrow A$$

be an observable.

For a cost bound $b \in B$, define

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

This is the set of **attainable observations below** b .

Equivalently, it is the set of observable values attained by objects admitting a Γ -description of cost at most b .

It may also be called the **bounded observable image** of $\mathcal{X}_{\Gamma}^{\leq b}$ under ω .

If $b \leq b'$, then

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

Thus

$$b \mapsto \text{Val}_{\Gamma, \omega}(b)$$

is an increasing profile of observable values.

The phrase **attainable observations below** b will be used for this set.

21 Observable profiles

The **observable profile** of ω relative to Γ is the function

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

It records how the observable sees the bounded-presentation filtration.

Instead of the full set $\text{Val}_{\Gamma, \omega}(b)$, one may study numerical summaries, such as

$$b \mapsto |\text{Val}_{\Gamma, \omega}(b)|,$$

$$b \mapsto \dim \text{span}(\text{Val}_{\Gamma, \omega}(b)),$$

or other sizes, ranks, volumes, supports, or distributional statistics.

However, the full observable profile is the primary object.

22 Observable lower-bound principle

Let

$$x \in \mathcal{X}.$$

If

$$\omega(x) \notin \text{Val}_{\Gamma, \omega}(b),$$

then

$$x \notin \mathcal{X}_{\Gamma}^{\leq b}.$$

Equivalently, x has no Γ -description of cost at most b .

When presentation complexity is numerical in an ordered resource scale, this gives

$$C_{\Gamma}(x) \not\leq b.$$

In a totally ordered numerical scale this is the familiar statement $C_{\Gamma}(x) > b$.

This is the basic observable lower-bound principle.

23 Cost-observable bounds

In practice, one often works with proved outer bounds for the full set

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

One proves an inclusion

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

for some explicit subset

$$S_b \subseteq A.$$

Then

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

implies

$$x \notin \mathcal{X}_{\Gamma}^{\leq b}.$$

Thus

$$\omega(x) \notin S_b \implies C_{\Gamma}(x) \not\leq b$$

whenever numerical presentation complexity is defined.

For totally ordered numerical costs this may be written as $C_{\Gamma}(x) > b$.

This gives the usual form of a lower-bound criterion:

low-cost descriptions force observations in S_b ;

the object has observation outside S_b ;

therefore the object lies outside the bounded presentation class.

24 Observed sizes

Suppose the observation space A has a size function

$$s : A \rightarrow R$$

to another ordered resource scale (R) . Then

$$s \circ \omega : \mathcal{X} \rightarrow R$$

is an **observed size**.

It is a numerical or resource-valued invariant of objects.

An observed size becomes a presentation complexity when it is induced by a presentation system.

The important mathematical question is often to relate observed sizes to presentation complexity. For example, if one proves

$$x \in \mathcal{X}_\Gamma^{\leq b} \implies s(\omega(x)) \leq F(b),$$

then

$$s(\omega(x)) > F(b) \implies x \notin \mathcal{X}_\Gamma^{\leq b}.$$

When numerical presentation complexity is defined,

$$s(\omega(x)) > F(b) \implies C_\Gamma(x) \not\leq b.$$

This covers many familiar lower-bound arguments: degree too large, rank too large, denominator too large, dimension too large, relation lattice too constrained, or observable distribution outside the attainable range.

A useful distinction is:

presentation costs come from description systems,
observed sizes come from observables.

25 Families of observables

A **family of observables** on $\mathcal{X}$ is a collection

$$\Omega = \omega_i : \mathcal{X} \rightarrow A_{ii \in I}.$$

It may be viewed as a joint observable

$$\omega_\Omega : \mathcal{X} \rightarrow \prod_{i \in I} A_i$$

defined by

$$\omega_\Omega(x) = (\omega_i(x))_{i \in I}.$$

The family Ω is **complete on a subclass** $\mathcal{C} \subseteq \mathcal{X}$ if, for all $x, y \in \mathcal{C}$,

$$\omega_i(x) = \omega_i(y) \quad \text{for every } i \in I$$

implies

$$x = y,$$

or, in an isomorphism-based context,

$$x \cong y.$$

A central problem is to determine which families of observables are complete on which bounded parts or natural subclasses.

26 Observable refinement, coarsening, and products

Observables carry a simple algebra.

If

$$\omega : \mathcal{X} \rightarrow A$$

is an observable and

$$q : A \rightarrow B$$

is a map, then

$$q \circ \omega : \mathcal{X} \rightarrow B$$

is a **coarsening** of ω .

An observable

$$\eta : \mathcal{X} \rightarrow E$$

refines ω if $\omega = q \circ \eta$ for some map $q : E \rightarrow A$. Refinement can only distinguish at least as much as the original observable.

Given observables

$$\omega : \mathcal{X} \rightarrow A \quad \text{and} \quad \eta : \mathcal{X} \rightarrow E,$$

their product observable is

$$(\omega, \eta) : \mathcal{X} \rightarrow A \times E.$$

For any presentation system Γ and cost bound b ,

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

The inclusion can be strict. Strictness records correlations among observations imposed by the common bounded presentation class.

27 Indistinguishability fibres

Given an observable

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

the fibre

$$\omega^{-1}(a)$$

is the class of objects indistinguishable by ω from the observable value (a).

If

$$\omega(x) = \omega(y),$$

then any method depending on ω gives the same output on x and y .

Thus every observable carries an intrinsic equivalence relation: it distinguishes objects up to its fibres.

For a family of observables Ω , define

$$x \sim_{\Omega} y$$

if

$$\omega(x) = \omega(y)$$

for every $\omega \in \Omega$.

The equivalence classes of $\sim_{\Omega}$ are the **indistinguishability fibres** of Ω .

Large fibres indicate blindness; small or trivial fibres indicate completeness.

Studying these fibres is essential when observables are used for classification.

28 Reconstruction by observables

A family of observables may do more than distinguish objects. It may reconstruct a structural invariant.

Let

$$\tau : \mathcal{X} \rightarrow T$$

be a structural invariant, such as a module, tensor, rank profile, support profile, lattice, isomorphism class, or moduli parameter.

A family Ω **reconstructs** τ on $\mathcal{C} \subseteq \mathcal{X}$ if

$$x \sim_{\Omega} y \implies \tau(x) = \tau(y)$$

for all $x, y \in \mathcal{C}$.

Equivalently, τ factors through the observational quotient

$$\mathcal{C} / \sim_{\Omega} .$$

If τ is a complete invariant on $\mathcal{C}$, then reconstruction of τ implies classification.

Many concrete results have this form:

word distributions reconstruct a rank profile,

commutator moments reconstruct a tensor pencil,

hom-count profiles reconstruct a module,

finite differences reconstruct coefficient data.

Thus reconstruction is a central use of observables beyond lower bounds.

29 Costed observables

Observations may also have costs.

A **costed observable family** on $\mathcal{X}$ is a triple

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

where:

$$\Omega$$

is a class of observables,

$$R$$

is an observation-resource scale, and

$$\chi : \Omega \rightarrow R$$

assigns an observation cost to each observable.

For an observation bound $r \in R$, define

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

Thus $\Omega^{\leq r}$ is the family of observables available with observation budget r .

Examples of observation costs include arity of a word observable, length of a word, number or order of test groups, degree of a polynomial probe, rank or matrix size of a linear observable, number of evaluations, depth of a logical formula, or quantifier rank.

A declared costed observable family is also an **observable scheme**: it specifies which tests are allowed before a lower-bound argument begins. If an observable is selected using target-dependent information, then the selection procedure must be canonical or its cost must be included in χ . Otherwise singleton tests would trivialize distinguishability; such tests are part of the theory only when the declared observable scheme explicitly allows and charges for them.

This distinction is essential:

presentation cost of objects

and

observation cost of tests

are different resources.

Costed observables form a second resource layer. They measure the complexity of observing, testing, and distinguishing once the object class and presentation systems have been fixed.

30 Distinguishing cost

Let

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

be a costed observable family.

For two objects $x, y \in \mathcal{X}$, define their **distinguishing cost** by

$$\text{dc}_{\mathcal{D}}(x, y) = \inf\{\chi(\omega) : \omega \in \Omega, \omega(x) \neq \omega(y)\}.$$

If no observable in Ω distinguishes x and y , then

$$\text{dc}_{\mathcal{D}}(x, y) = \infty.$$

This measures the minimum observation budget needed to tell x and y apart using the allowed observables.

Observable-barrier results have the form:

$$\Omega^{\leq r} \text{ is not complete on } \mathcal{C},$$

but

$$\Omega^{\leq r'} \text{ is complete on } \mathcal{C}$$

for some larger r' , often with r' optimal.

Such results show that observation itself has irreducible resource requirements.

31 Normal-form compilers

Many applications depend on a theorem that converts raw descriptions or raw observables into a simpler algebraic normal form.

A **normal-form compiler** is a result showing that a class of descriptions or observables is equivalent to a more structured class.

For descriptions, a normal-form compiler may have the form:

$$d \in \mathcal{D}^{\leq b} \implies \exists d' \in \mathcal{D}_{\text{nf}}^{\leq f(b)} \text{ with } \rho(d') = \rho(d).$$

This says that every low-cost description can be replaced by a normal-form description with controlled overhead.

For observables, a normal-form compiler may have the form

$$\omega_{\text{raw}} \rightsquigarrow \omega_{\text{nf}},$$

where the raw observable and the normal-form observable determine the same information on the class under study.

Examples include:

word equations in class-two groups $\rightsquigarrow$ linear and alternating quadratic data,

relative word observables $\rightsquigarrow$ matrix-rank observables,

Boolean lift constraints $\rightsquigarrow$ Möbius coefficient constraints,

commutator-word moments $\rightsquigarrow$ rank probes of alternating tensors,

coordinate descriptions $\rightsquigarrow$ primitive height-normalized coordinates.

Normal-form compilers are often the technical heart of applications. They make bounded objects and attainable observations computable, constrainable, or reconstructive.

32 Complete normal forms as resolutions

A stronger object is a complete normal form for the realization map itself.

Let

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

be a realization map. A **complete normal form** for ρ is a computable map

$$N : \mathcal{D} \rightarrow \mathcal{N}$$

to an effectively coded space of normal forms with decidable equality, such that

$$N(d) = N(e) \iff \rho(d) \equiv_{\mathcal{X}} \rho(e).$$

Such an N is a computable resolution of the equivalence relation induced by ρ . It replaces the problem of deciding whether two descriptions realize the same object by equality of their normal forms.

This is stronger than a cost-controlled compiler into a convenient class of descriptions. A compiler may give good representatives in a bounded range or for a structured subclass. A complete normal form resolves every fibre at once.

Consequently, undecidable fibres obstruct complete computable normalization: if d_x is a description of x , then a complete normal form decides membership in $\mathfrak{F}_\rho(x)$ by testing whether $N(d) = N(d_x)$.

33 Realization images and algebraic obstructions

A presentation system contains a realization map

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

Often the key problem is to understand the image of ρ , or the image of a bounded, evaluated, linearized, or reduced version of ρ .

For a bound b , the bounded image is

$$\mathcal{X}_\Gamma^{\leq b} = \rho(\mathcal{D}^{\leq b}).$$

In algebraic settings, obstructions may arise from failure to lie in an image, cokernel classes, congruence conditions, rank constraints, lattice constraints, support constraints, or coefficient alphabet restrictions.

Thus some lower bounds are not initially about size. They are about realizability.

The general pattern is:

low-cost descriptions land inside a constrained image;

the target object or observation lies outside that image;

therefore

no low-cost description exists.

This is especially important when the presentation system is algebraic, linear, modular, or arithmetic.

34 Pullback and transport of observables

Let

$$F : \mathcal{X} \rightarrow \mathcal{Y}$$

be a map of object classes, and let

$$\eta : \mathcal{Y} \rightarrow A$$

be an observable on $\mathcal{Y}$.

Then

$$\omega = \eta \circ F : \mathcal{X} \rightarrow A$$

is an observable on $\mathcal{X}$.

This is the **pullback observable** induced by F and η .

If F is cost-controlled with monotone overhead f , then

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

Therefore,

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

Thus maps between object classes transport observable profiles.

This is useful when $(F(x))$ is a quotient, reduction, abelianization, associated graded object, local model, or simpler invariant of x .

A transported lower bound has the following form.

If

$$C_{\Gamma_{\mathcal{Y}}}(F(x)) \not\leq f(b),$$

then

$$C_{\Gamma_{\mathcal{X}}}(x) \not\leq b.$$

Thus a lower bound for a simpler transported object can imply a lower bound for the original object.

35 Relative and fibre methods

Sometimes an object is studied through a map

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

The object $q(x)$ may be a quotient, reduction, base object, abelianization, associated graded object, invariant, or observable component of x .

The fibre

$$q^{-1}(q(x))$$

contains the information not determined by $q(x)$.

One may try to decompose the cost of x into the cost of the base object $q(x)$ and the residual cost of specifying x inside the fibre.

Informally,

$$\text{cost of } x \approx \text{cost of } q(x) + \text{cost of lifting data.}$$

This method is useful for extensions, quotients, fibrations, local-to-global problems, model-plus-residual descriptions, cohomological classification, and situations where an observable sees only part of the object.

A full relative theory may introduce conditional presentation systems or relative costs

$$C_\Gamma(x \mid q(x)).$$

But the core idea is already present in the basic formalism through fibres of maps and observables.

36 Bounded-counterexample arguments

Let P be a property of objects in $\mathcal{X}$. Define the bad set

$$B_P = \{x \in \mathcal{X} : P(x) \text{ fails}\}.$$

Thus proving P for all objects in $\mathcal{X}$ is the same as proving

$$B_P = \emptyset.$$

Let Γ be a presentation system on $\mathcal{X}$, and let b be a cost bound.

A **bounded-counterexample principle** for P , with respect to Γ and b , is an implication of the form

$$B_P \neq \emptyset \implies B_P \cap \mathcal{X}_\Gamma^{\leq b} \neq \emptyset.$$

In words: if P fails somewhere, then it already fails on a Γ -presentable object of cost at most b .

This principle reduces a universal statement to the bounded part $\mathcal{X}_\Gamma^{\leq b}$: any failure of P has a low-cost counterexample.

To prove P , one combines it with an exclusion result:

$$B_P \cap \mathcal{X}_\Gamma^{\leq b} = \emptyset.$$

Together, the two statements imply

$$B_P = \emptyset.$$

Indeed, if $B_P \neq \emptyset$, the bounded-counterexample principle gives an element of

$$B_P \cap \mathcal{X}_\Gamma^{\leq b},$$

contradicting the exclusion result.

The exclusion result may be proved directly, by finite enumeration, by a normal-form theorem, or by observables.

The observable form is the following. Suppose there is an observable

$$\omega : \mathcal{X} \rightarrow A$$

and a subset $S_b \subseteq A$ such that all objects of cost at most b have observable value in S_b :

$$\mathcal{X}_\Gamma^{\leq b} \subseteq \omega^{-1}(S_b),$$

while every bad object has observable value outside S_b :

$$B_P \subseteq \mathcal{X} \setminus \omega^{-1}(S_b).$$

Then no bad object can lie in the bounded part. If bad objects, if they exist, must lie in the bounded part, then no bad objects exist.

In slogan form:

counterexamples must be cheap,
the bounded class contains no counterexamples,
therefore no counterexamples exist.

The argument has a precise form: an independent mathematical reason forces counterexamples into a bounded presentation class, and a second argument excludes counterexamples from that bounded class.

Proposition 36.1 (No universal bounded-counterexample procedure). *There is no algorithm which, given an arbitrary decidable predicate $P(n)$ on $\mathbb{N}$, outputs a finite set $S \subseteq \mathbb{N}$ such that*

$$\exists n \neg P(n) \implies \exists s \in S \neg P(s).$$

Proof. If such an algorithm existed, it would decide the halting problem. Given a Turing machine M , let $P_M(t)$ be the decidable predicate saying that M has not halted within t steps. Apply the hypothetical algorithm to P_M , obtaining a finite set S . Checking $P_M(s)$ for all $s \in S$ would decide whether some counterexample to P_M exists, equivalently whether M ever halts. $\square$

Thus bounded-counterexample principles require separate mathematical input. They are powerful when available, but they cannot be generated uniformly from decidability alone.

37 Reusable General Theorems

The formalism gives stable proof patterns. The following theorems are abstract templates for height balls, coefficient boxes, word observables, hom-count tests, rank probes, normal forms, proof languages, and bounded counterexample searches.

Theorem 37.1 (Bounded realization-image obstruction). *Let $\Gamma = (\mathcal{D}, \rho, \kappa)$ be a presentation system on $\mathcal{X}$. Let*

$$\theta : \mathcal{X} \rightarrow T$$

be any invariant or observable, and suppose that for a cost bound b one has a proved inclusion

$$\theta(\mathcal{X}_\Gamma^{\leq b}) \subseteq S_b$$

for some subset $S_b \subseteq T$. Then

$$\theta(x) \notin S_b \implies x \notin \mathcal{X}_\Gamma^{\leq b}.$$

If numerical presentation complexity is defined, then $C_\Gamma(x) \not\leq b$. 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 the proved inclusion $\theta(x) \in S_b$. The contrapositive gives the claim. $\square$

This is the common lower-bound pattern behind coefficient constraints, bounded realization images, rank constraints, height denominators, volume estimates, and algebraic support restrictions. The bounded set $\mathcal{X}_{\Gamma}^{\leq b}$ supplies the constrained image S_b .

Theorem 37.2 (Normal-form transfer). *Let Γ_{raw} and Γ_{nf} be presentation systems for the same object class $\mathcal{X}$, with*

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

and

$$\Gamma_{\text{nf}} = (\mathcal{D}_{\text{nf}}, \rho_{\text{nf}}, \kappa_{\text{nf}}).$$

Suppose there is a compiler

$$T : \mathcal{D}_{\text{raw}} \rightarrow \mathcal{D}_{\text{nf}}$$

and a monotone overhead function f such that

$$\rho_{\text{nf}}(T(d)) = \rho_{\text{raw}}(d)$$

and

$$\kappa_{\text{nf}}(T(d)) \leq f(\kappa_{\text{raw}}(d))$$

for all $d \in \mathcal{D}_{\text{raw}}$. Then

$$\mathcal{X}_{\Gamma_{\text{raw}}}^{\leq b} \subseteq \mathcal{X}_{\Gamma_{\text{nf}}}^{\leq f(b)}.$$

Consequently, every lower bound against Γ_{nf} -normal forms below $f(b)$ is also a lower bound against raw presentations below b .

Proof. If $x = \rho_{\text{raw}}(d)$ with $\kappa_{\text{raw}}(d) \leq b$, then $T(d)$ is a normal-form description of the same object and has cost at most $f(b)$. $\square$

The compiler is the mathematical content: the cost inequality turns a normal form into a transfer theorem.

Theorem 37.3 (Observable factorization). *Let $\mathfrak{D} = (\Omega, \chi, R)$ be a costed observable family on $\mathcal{X}$, and let $r \in R$. Let $\mathcal{C} \subseteq \mathcal{X}$, and suppose there is a structural statistic*

$$\tau_r : \mathcal{C} \rightarrow T_r$$

such that every observable of cost at most r factors through τ_r on $\mathcal{C}$:

$$\omega = a_{\omega} \circ \tau_r \quad (\omega \in \Omega^{\leq r})$$

for suitable maps $a_{\omega} : T_r \rightarrow A_{\omega}$. Then

$$\tau_r(x) = \tau_r(y) \implies \omega(x) = \omega(y) \quad \text{for all } \omega \in \Omega^{\leq r}.$$

Thus observables of budget r distinguish objects only up to the fibres of τ_r .

Proof. If $\tau_r(x) = \tau_r(y)$, then $a_{\omega}(\tau_r(x)) = a_{\omega}(\tau_r(y))$ for every factorizing observable ω . $\square$

This theorem is the basic mechanism of observable complexity: bounded observables see the compressed statistic τ_r .

Corollary 37.4 (Finite-observable blindness). *In the setting of the observable factorization theorem, if one fibre of τ_r contains two non-equivalent objects $x, y \in \mathcal{C}$, then $\Omega^{\leq r}$ is not complete on $\mathcal{C}$. In particular,*

$$\text{dc}_{\mathfrak{D}}(x, y) \not\leq r,$$

provided the distinguishing cost is interpreted in the preorder R .

This is the abstract form of many bounded-observable blindness results.

Theorem 37.5 (Reconstruction from factored observables). *Keep the notation of the observable factorization theorem. Suppose, conversely, that τ_r is determined by the observables $\Omega^{\leq r}$: whenever $x, y \in \mathcal{C}$ satisfy*

$$\omega(x) = \omega(y) \quad (\omega \in \Omega^{\leq r}),$$

one has $\tau_r(x) = \tau_r(y)$. Then the observational quotient of $\mathcal{C}$ by $\Omega^{\leq r}$ is exactly the quotient by the fibres of τ_r . If τ_r is a complete invariant on $\mathcal{C}$, then $\Omega^{\leq r}$ is complete on $\mathcal{C}$.

Proof. The observable factorization theorem gives one implication: equal τ_r implies equal observations. The hypothesis gives the reverse implication. Hence the two equivalence relations coincide. $\square$

Proposition 37.6 (Cost-sensitive fibration bound). *Let $q : \mathcal{X} \rightarrow \mathcal{Y}$ be a structural map. Suppose $\mathcal{Y}$ has a presentation system $\Gamma_{\mathcal{Y}}$, and suppose that for every $y \in \mathcal{Y}$ there is a relative presentation system $\Gamma_{\text{rel}}(y)$ for the fibre $q^{-1}(y)$. Assume objects of $\mathcal{X}$ can be assembled from base and relative descriptions with overhead A :*

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

Then

$$C_{\Gamma_{\mathcal{X}}}(x) \leq \inf_{e, h} A(\kappa_{\mathcal{Y}}(e), \kappa_{\text{rel}}(h)),$$

where the infimum ranges over descriptions e of $q(x)$ and relative descriptions h of x in the fibre over $q(x)$.

Proof. Every admissible pair (e, h) assembles to a description of x with the stated cost. Taking the infimum gives the bound. $\square$

This proposition is an upper-bound assembly statement. Relative lower bounds arise from converse hypotheses showing that every low-cost description of x yields low-cost base and fibre data.

Theorem 37.7 (Unbounded verification-length principle). *Let $S \subseteq \{0, 1\}^*$ be recursively enumerable but not decidable. Suppose there is a decidable relation $R(u, \pi)$ such that*

$$u \in S \iff \exists \pi R(u, \pi).$$

For $u \in S$, define

$$m_R(u) = \min\{|\pi| : R(u, \pi)\},$$

and define

$$b_R(n) = \max_{\substack{|u| \leq n \\ u \in S}} m_R(u),$$

with value 0 if the maximum is taken over an empty set. Then b_R is not bounded above by any total computable function.

Proof. Suppose $b_R(n) \leq F(n)$ for some total computable function F . Given u , compute $F(|u|)$ and check the finitely many words π with $|\pi| \leq F(|u|)$. Since R is decidable, this finite search decides whether $u \in S$. This contradicts the undecidability of S . $\square$

The theorem is elementary but important: recursively enumerable membership, once undecidable, forces shortest membership verification data to escape every computable uniform bound.

Theorem 37.8 (Undecidable fibres force nonrecursive verification-length growth). *Let $\rho : \mathcal{D} \rightarrow \mathcal{X}$ be a realization map with effectively coded descriptions and coding length ℓ . Fix $x \in \mathcal{X}$, and suppose the fibre*

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

is recursively enumerable but not decidable. Suppose moreover that membership in this fibre is represented by a decidable verification relation

$$\rho(d) = x \iff \exists \pi R_x(d, \pi).$$

Then the bounded verification-length profile

$$L_{R_x}(n) = \max_{\substack{\ell(d) \leq n \\ d \in \mathfrak{F}_\rho(x)}} \min\{|\pi| : R_x(d, \pi)\}$$

is not bounded above by any total computable function.

Proof. Apply the unbounded verification-length principle to the recursively enumerable undecidable set $\mathfrak{F}_\rho(x)$, using the effective coding of descriptions. $\square$

This turns undecidability of a presentation fibre into a quantitative lower bound on membership verification.

Corollary 37.9 (Effective reversible move systems force nonrecursive bridge growth). *Let $\mathcal{M}$ be an effective finite-branching reversible move system on $\mathcal{D}$ preserving realization. Fix $d_x \in \mathfrak{F}_\rho(x)$. Suppose*

$$d \in \mathfrak{F}_\rho(x) \iff d \text{ is connected to } d_x \text{ by an } \mathcal{M}\text{-path.}$$

If membership in $\mathfrak{F}_\rho(x)$ is undecidable, then the bridge profile

$$B_{\rho, \mathcal{M}, x, d_x}(n) = \sup_{\substack{\ell(d) \leq n \\ d \in \mathfrak{F}_\rho(x)}} d_{\mathcal{M}}(d, d_x)$$

is not bounded above by any total computable function.

Proof. If $B_{\rho, \mathcal{M}, x, d_x}(n) \leq F(n)$ for a total computable F , then membership in the fibre could be decided as follows. Given d , search the finite ball of $\mathcal{M}$ -paths of length at most $F(\ell(d))$ starting at d , and check whether d_x occurs. By the connectivity hypothesis this decides $d \in \mathfrak{F}_\rho(x)$, a contradiction. $\square$

Theorem 37.10 (No complete computable resolution). *Let $\rho : \mathcal{D} \rightarrow \mathcal{X}$ be a realization map with effectively coded descriptions. Suppose there is a computable map*

$$N : \mathcal{D} \rightarrow \mathcal{N}$$

to an effectively coded set $\mathcal{N}$ with decidable equality such that

$$N(d) = N(e) \iff \rho(d) \equiv_{\mathcal{X}} \rho(e).$$

Then every represented fibre $\mathfrak{F}_\rho(x)$, with a fixed description d_x satisfying $\rho(d_x) = x$, is decidable.

Proof. Given d , compute $N(d)$ and $N(d_x)$. Then

$$d \in \mathfrak{F}_\rho(x) \iff N(d) = N(d_x).$$

Since equality in $\mathcal{N}$ is decidable, the fibre membership problem is decidable. $\square$

Consequently, if a represented fibre is undecidable, the realization map admits no complete computable normal form in the above sense. Normal forms of this strength are resolutions of the fibre equivalence relation.

Corollary 37.11 (Non-collapse of fibre complexity). *There are presentation systems with an object x satisfying*

$$C_\Gamma(x) = 0$$

while the verification-length profile of $\mathfrak{F}_\rho(x)$ is not bounded above by any total computable function. Hence no function of minimal presentation complexity alone can bound fibre recognition or membership-verification cost in general.

Proof. Let $S \subseteq \{0, 1\}^*$ be recursively enumerable but undecidable, with decidable verification relation $R(u, \pi)$. Let

$$\mathcal{D} = \{*\} \sqcup \{d_u : u \in \{0, 1\}^*\}.$$

Let $\equiv$ be the equivalence relation generated by

$$d_u \equiv * \iff u \in S,$$

and let $\mathcal{X} = \mathcal{D}/\equiv$, with $\rho : \mathcal{D} \rightarrow \mathcal{X}$ the quotient map. Put $\kappa(*) = 0$ and $\kappa(d_u) = 1 + |u|$. For $x = \rho(*)$, one has $C_\Gamma(x) = 0$, while $d_u \in \mathfrak{F}_\rho(x)$ exactly when $u \in S$. The unbounded verification-length principle gives the claimed nonrecursive growth. $\square$

This construction is deliberately schematic. Natural non-collapse examples require specific undecidability or lower-bound results.

The point of the corollary is that minimal presentation complexity alone cannot be a general bound for fibre recognition or membership-verification growth.

Theorem 37.12 (Prefix-free search penalty). *Let (X, μ) be a probability space, and let $\{A_c\}_{c \in \mathcal{C}}$ be a countable family of searched events. Suppose*

$$\mu(A_c) \leq p_c$$

with $p_c > 0$, and suppose a search cost $K(c)$ satisfies Kraft's inequality

$$\sum_{c \in \mathcal{C}} 2^{-K(c)} \leq 1.$$

Then, for every $t \geq 0$,

$$\mu \left(\bigcup_{\{c: -\log_2 p_c - K(c) \geq t\}} A_c \right) \leq 2^{-t}.$$

Proof. By the union bound,

$$\mu \left(\bigcup_{\{c: -\log_2 p_c - K(c) \geq t\}} A_c \right) \leq \sum_{\{c: -\log_2 p_c - K(c) \geq t\}} p_c.$$

For every index in the displayed sum,

$$p_c \leq 2^{-t} 2^{-K(c)}.$$

Hence

$$\sum_{\{c: -\log_2 p_c - K(c) \geq t\}} p_c \leq 2^{-t} \sum_c 2^{-K(c)} \leq 2^{-t}.$$

□

This is the formal search-penalty principle behind adaptive discovery. The relevant budget is the description cost of the searched event. For N independent samples, one replaces p_c by p_c^N , producing the score $N \log_2(1/p_c) - K(c)$.

Theorem 37.13 (Cheap-counterexample exclusion). *Let P be a property of objects in $\mathcal{X}$, and let*

$$B_P = \{x \in \mathcal{X} : P(x) \text{ fails}\}.$$

Suppose a presentation system Γ and a bound b satisfy:

$$B_P \neq \emptyset \implies B_P \cap \mathcal{X}_\Gamma^{\leq b} \neq \emptyset$$

and

$$B_P \cap \mathcal{X}_\Gamma^{\leq b} = \emptyset.$$

Then P holds for all objects in $\mathcal{X}$.

Proof. If B_P were nonempty, the first hypothesis would produce an element of $B_P \cap \mathcal{X}_\Gamma^{\leq b}$, contradicting the second hypothesis. □

This theorem captures finite-reduction arguments: counterexamples must be cheap, but cheap counterexamples are excluded.

38 Basic examples

38.1 Heights

For rational points in projective space,

$$\mathcal{X} = \mathbb{P}^n(\mathbb{Q}),$$

one may take descriptions to be primitive integer coordinate vectors

$$(a_0, \dots, a_n) \in \mathbb{Z}_{\text{prim}}^{n+1} \setminus \{0\},$$

with realization

$$\rho(a_0, \dots, a_n) = [a_0 : \dots : a_n]$$

and cost

$$\kappa(a_0, \dots, a_n) = \log \max_i |a_i|.$$

The resulting presentation complexity is the logarithmic height.

Thus height theory is an arithmetic instance of presentation complexity.

Height balls are bounded parts. Height inequalities under morphisms are cost-controlled maps. Counting points of bounded height is volume theory for bounded parts.

38.2 Kolmogorov complexity

For finite binary strings,

$$\mathcal{X} = \{0, 1\}^*,$$

take descriptions to be programs for a universal machine U :

$$\mathcal{D} = \{0, 1\}^*,$$

with realization

$$\rho(p) = U(p)$$

and cost

$$\kappa(p) = |p|.$$

Then

$$K_U(x) = \min\{|p| : U(p) = x\}$$

is Kolmogorov complexity relative to U .

38.3 Circuit complexity

For Boolean functions,

$$\mathcal{X} = \{f : \{0, 1\}^n \rightarrow \{0, 1\}\},$$

descriptions may be Boolean circuits, realization is the computed function, and cost is circuit size or depth.

The presentation complexity is minimum circuit size or depth.

38.4 Group presentations

For finite groups, one possible presentation system uses finite presentations

$$\langle S \mid R \rangle,$$

with realization

$$\rho(S, R) = F(S) / \langle\langle R \rangle\rangle$$

and cost such as

$$\kappa(S, R) = (|S|, |R|, \sum_{r \in R} |r|).$$

Another presentation system may use multiplication tables, word circuits, semidirect-product data, module data, commutator tensors, finite local models, or permutation-cover data.

Different presentation systems measure different mathematical resources.

Words may play two distinct roles. As relators, words are part of descriptions and contribute to presentation cost. As word maps or word equations, words define observables and contribute to observation cost.

38.5 Algebraic and straight-line complexity

For polynomials or algebraic expressions, descriptions may be arithmetic circuits or straight-line programs.

The realization map sends a circuit to the polynomial or function it computes.

The cost may be the number of operations, depth, degree of allowed gates, or circuit size.

Observables may include degree, rank of partial derivative spaces, support, evaluation data, or factorization patterns.

This gives a presentation-theoretic view of algebraic complexity.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, ChatGPT, by OpenAI, was used to assist with mathematical drafting, formalization, review, and editing. This work is shared as a preliminary AI-assisted mathematical note. The mathematical content may have been only partially reviewed and may contain errors; it should not be treated as peer-reviewed or as a fully verified manuscript.

References

- [1] S. Arora and B. Barak, *Computational Complexity: A Modern Approach*, Cambridge University Press, 2009.
- [2] A. N. Kolmogorov, Three approaches to the quantitative definition of information, *Problems of Information Transmission* 1 (1965), 1–7.
- [3] M. Li and P. Vitanyi, *An Introduction to Kolmogorov Complexity and Its Applications*, third edition, Springer, 2008.
- [4] D. G. Northcott, An inequality in the theory of arithmetic on algebraic varieties, *Proceedings of the Cambridge Philosophical Society* 45 (1950), 502–509.