

Presentation Theory III: Normal-Form Compilation, Residual Fibres, and Stable Canonical Access

Luca Blanchi

June 2026

Abstract

This note develops a normal-form layer for Presentation Theory. A normal form is treated not as a privileged representative in isolation, but as a costed access transformation: it converts raw descriptions into more rigid data while controlling overhead, verification data, residual fibres, and the observables that survive the transformation. This point of view gives a calculus of normal-form packages, their composition, observable quotients, residual decompositions, and lower bounds from separation of budgeted normal images.

The formal results are deliberately elementary but reusable. Bounded compilation sends bounded raw objects into bounded normal images. Separators of those images give lower bounds. Observable families and residual equivalence relations form a Galois connection. Residual complexity splits the cost of reconstructing an object into normal data and fibre data. Fragmented towers may be complete only in a non-effective sense, and their optimal truncation functions have the same Turing degree as the underlying positive predicate. In filtered effective systems, computable minimal selectors exist exactly when bounded presentation is decidable, while computable canonical selectors exist exactly when both bounded presentation and equivalence are decidable. Heighted normal forms satisfy no-free-normal-form inequalities: bounded normal data and bounded gauges cannot produce objects of unbounded height, and degenerating normal forms must pay for the path producing the limit. Rewriting systems supply normal forms when termination, local confluence, and cost growth are controlled.

The main concrete application is a fixed-gap instability theorem for spectral canonicalization. On the orbit of real symmetric matrices with fixed simple spectrum and uniformly positive eigengap, no global eigenbasis canonicalizer can be continuous. In fact any deterministic sign choice has jumps of size approaching 2. The stable replacement is the projector-valued, or stacky, spectral normal form. This explains, in normal-form language, why sign- and basis-invariant spectral graph methods are structurally natural: continuous basis-invariant spectral encodings factor through the relevant projector data.

1 Purpose

The first note on Presentation Theory introduced presentation systems, realization fibres, observables, verification costs, and basic bounded-image principles. The second note developed controlled transfer packages: maps that transport access to objects rather than just objects themselves. This note isolates a third recurring mechanism: normal-form compilation.

Classical normal forms include row echelon form, Smith normal form, reduced words, reduced Grobner bases, Jordan form, Chomsky normal form, and canonical graph labels. These examples suggest a common operation:

$$\text{raw description} \longmapsto \text{rigid datum}.$$

For Presentation Theory this is not enough. A useful normal form should also say what it costs, what it verifies, which observables it preserves, what information remains in the residual fibre, and which lower bounds become visible after normalization.

The guiding principle is:

a normal form is controlled rigidification of access.

This note keeps the framework intentionally modest. It does not claim that every mathematical classification problem has a natural normal form, nor that a universal obstruction method exists. It gives a language for tracking what a normal-form method actually controls.

2 Presentation Contexts and Resource Scales

Definition 2.1 (Resource scale). *A resource scale is a preordered commutative monoid*

$$(B, \leq, +, 0).$$

Its elements are budgets. Typical examples are

$$\mathbb{N}, \quad \mathbb{N}^r, \quad \mathbb{R}_{\geq 0}, \quad \mathbb{N} \times \mathbb{R}_{\geq 0}.$$

Maps between resource scales are assumed monotone unless stated otherwise.

Definition 2.2 (Presentation system). *A presentation system is a triple*

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

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

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

is a realization map into an object class $\mathcal{X}$, and

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

is a cost function.

For $b \in B$, define

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

and

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

When the infimum is meaningful, the presentation complexity of $x \in \mathcal{X}$ is

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

The realization fibre of x is

$$\mathfrak{F}_\rho(x) = \rho^{-1}(x).$$

Definition 2.3 (Presentation context). *A presentation context is a tuple*

$$\mathfrak{P} = (\Gamma, \Omega, \mathcal{L}),$$

where Γ is a presentation system, Ω is a chosen family of observables on the realized object class, and $\mathcal{L}$ is a language of properties, tests, or predicates interpreted through the presentation.

Observables may also be costed. We write $\Omega^{\leq r}$ for observables of budget at most r . Their common fibres record the information invisible at that observational budget.

3 Normal-Form Packages

Definition 3.1 (Normal-form package). *Let*

$$\Gamma_0 = (\mathcal{D}_0, \rho_0, \kappa_0), \quad \Gamma_1 = (\mathcal{D}_1, \rho_1, \kappa_1)$$

be presentation systems. A normal-form package

$$\mathcal{N} : \Gamma_0 \rightsquigarrow \Gamma_1$$

consists of:

(i) *a compilation relation*

$$\text{Comp}_{\mathcal{N}}(d, n, \nu),$$

where $d \in \mathcal{D}_0$, $n \in \mathcal{D}_1$, and ν is verification data;

(ii) *a verification procedure*

$$\text{Check}_{\mathcal{N}}(d, n, \nu) \in \{0, 1\};$$

(iii) *comparison maps*

$$\Phi_{\mathcal{N}} : \mathcal{X}_0 \rightarrow \mathcal{Y}, \quad \Theta_{\mathcal{N}} : \mathcal{X}_1 \rightarrow \mathcal{Y};$$

(iv) *monotone overhead functions controlling normal size and verification cost;*

(v) *optionally, residual extraction and assembly data.*

The package is sound if

$$\text{Comp}_{\mathcal{N}}(d, n, \nu) \implies \Theta_{\mathcal{N}}(\rho_1(n)) = \Phi_{\mathcal{N}}(\rho_0(d)).$$

It is verifiably sound if

$$\text{Check}_{\mathcal{N}}(d, n, \nu) = 1 \implies \Theta_{\mathcal{N}}(\rho_1(n)) = \Phi_{\mathcal{N}}(\rho_0(d)).$$

Definition 3.2 (Completeness up to overhead). *A normal-form package is complete up to overhead $f_{\mathcal{N}}$ if, for every $d \in \mathcal{D}_0$ with $\kappa_0(d) \leq b$, there exist n and ν such that*

$$\text{Comp}_{\mathcal{N}}(d, n, \nu), \quad \kappa_1(n) \leq f_{\mathcal{N}}(b),$$

with verification data and verification time bounded by the declared overhead functions.

Definition 3.3 (Profiles). *The normal-size profile is*

$$F_{\mathcal{N}}(b) = \sup_{\kappa_0(d) \leq b} \inf \{ \kappa_1(n) : \exists \nu, \text{Comp}_{\mathcal{N}}(d, n, \nu) \}.$$

The verification-data profile is

$$V_{\mathcal{N}}(b) = \sup_{\kappa_0(d) \leq b} \inf \{ |\nu| : \exists n, \text{Comp}_{\mathcal{N}}(d, n, \nu) \}.$$

The verification-time profile is defined similarly. If residual data are present, one also records a residual profile

$$R_{\mathcal{N}}(b) = \sup_{\kappa_0(d) \leq b} C_{\text{rel}}(d \mid n),$$

where n ranges over normal data for d .

4 Modes of Normality

Different normal forms preserve different parts of an object.

Definition 4.1 (Representative normal form). *A representative normal form preserves the object itself:*

$$\rho_1(n) = \rho_0(d).$$

Definition 4.2 (Observable normal form). *An observable normal form preserves a declared family of observables. For each $\omega \in \Omega$, there is an induced observable a_ω on normal data such that*

$$\omega(\rho_0(d)) = a_\omega(n)$$

on compiled pairs.

Definition 4.3 (Relative normal form). *Given a map $q : \mathcal{X} \rightarrow \mathcal{Y}$, a relative normal form records a base datum $q(x)$ and leaves the remaining information in the fibre over $q(x)$.*

Definition 4.4 (Stratified and stacky normal forms). *A stratified normal form first assigns a stratum and then normalizes within that stratum. A stacky normal form records a datum only up to a groupoid of choices. In practice this means that a stable normal form may be a quotient datum together with a residual torsor or gauge group, rather than a single global representative.*

Definition 4.5 (Fragmented normal form). *A fragmented normal form is a tower*

$$\mathcal{N}_0, \mathcal{N}_1, \mathcal{N}_2, \dots$$

whose union is complete, while the required truncation level may be hard or impossible to bound effectively.

5 Basic Calculus

Theorem 5.1 (Bounded normal image). *Let*

$$\mathcal{N} : \Gamma_0 \rightsquigarrow \Gamma_1$$

be a representative normal-form package with overhead f . Then

$$\mathcal{X}_{\Gamma_0}^{\leq b} \subseteq \mathcal{X}_{\Gamma_1}^{\leq f(b)}.$$

More generally, if $\mathcal{N}$ is observable and $\omega \circ \rho_0 = a_\omega \circ \rho_1$ on compiled data, then

$$\omega(\mathcal{X}_{\Gamma_0}^{\leq b}) \subseteq a_\omega(\mathcal{X}_{\Gamma_1}^{\leq f(b)}).$$

Proof. If $x \in \mathcal{X}_{\Gamma_0}^{\leq b}$, then $x = \rho_0(d)$ for some d with $\kappa_0(d) \leq b$. By completeness, d compiles to n with $\kappa_1(n) \leq f(b)$. In the representative case, $\rho_1(n) = x$, so $x \in \mathcal{X}_{\Gamma_1}^{\leq f(b)}$. The observable statement follows by applying the factorization identity. $\square$

Corollary 5.2 (Normal-image obstruction). *Let $\theta : \mathcal{X}_1 \rightarrow Z$ be an invariant or observable on the normal side. Define*

$$S_{\mathcal{N}, \theta}(b) = \theta(\mathcal{X}_{\Gamma_1}^{\leq f(b)}).$$

If

$$\theta(x) \notin S_{\mathcal{N}, \theta}(b),$$

then $x \notin \mathcal{X}_{\Gamma_0}^{\leq b}$. Equivalently, $C_{\Gamma_0}(x) > b$.

Proof. This is the contrapositive of the bounded normal-image theorem. $\square$

Theorem 5.3 (Composition). *Suppose*

$$\Gamma_0 \xrightarrow{\mathcal{N}_{01}} \Gamma_1 \xrightarrow{\mathcal{N}_{12}} \Gamma_2$$

are normal-form packages with overheads f_{01} and f_{12} . Their composite has overhead

$$f_{02} = f_{12} \circ f_{01}.$$

Verification-data and verification-time profiles compose by concatenating the corresponding verification data and adding the compatibility cost.

Proof. For $\kappa_0(d) \leq b$, choose a compiled normal datum n_1 with $\kappa_1(n_1) \leq f_{01}(b)$. Apply the second package to n_1 , obtaining n_2 with

$$\kappa_2(n_2) \leq f_{12}(f_{01}(b)).$$

The verification data consist of the two verification records together with the data checking that the output of the first compilation is an admissible input for the second. $\square$

Remark 5.4 (Products, restrictions, and quotients). *Product normal forms combine several normal data. Restriction to a subclass may improve overhead. Quotient normal forms intentionally discard information, lowering cost at the price of enlarging residual fibres. These operations are often more useful than a search for a single complete representative.*

6 Observable Quotients and Residual Fibres

Let Ω be a family of observables on $\mathcal{X}$. Define

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

The quotient

$$Q_\Omega = \mathcal{X} / \sim_\Omega$$

is the universal normal datum preserving all observables in Ω .

Definition 6.1 (Residual relation). *For a family Ω , define*

$$\text{Res}(\Omega) = \{(x, y) : \omega(x) = \omega(y) \text{ for all } \omega \in \Omega\}.$$

For a relation $R \subseteq \mathcal{X} \times \mathcal{X}$, define

$$\text{Obs}(R) = \{\omega : \omega(x) = \omega(y) \text{ whenever } xRy\}.$$

Proposition 6.2 (Observable-residual Galois connection). *For every observable family Ω and relation R ,*

$$\Omega \subseteq \text{Obs}(R) \iff R \subseteq \text{Res}(\Omega).$$

Proof. The left-hand side says that every observable in Ω is constant on every R -pair. This is exactly the statement that every R -pair lies in the residual relation induced by Ω . $\square$

Definition 6.3 (Residual fibre). *Let $N : \mathcal{X} \rightarrow \mathcal{Y}$ be a normal datum map. The residual fibre over $y \in \mathcal{Y}$ is*

$$\text{Res}_N(y) = N^{-1}(y).$$

For $x \in \mathcal{X}$, write

$$\text{Res}_N(x) = \text{Res}_N(N(x)).$$

Theorem 6.4 (Residual decomposition). *Assume that a normal form has controlled extraction*

$$x \mapsto (N(x), r_x)$$

and controlled assembly

$$\text{Asm}(N(x), r_x) = x,$$

with assembly cost bounded by

$$\kappa_X(x) \leq A(\kappa_N(N(x)), \kappa_{\text{rel}}(r_x)).$$

Then

$$C_X(x) \leq A(C_N(N(x)), C_{\text{rel}}(x \mid N(x))).$$

Conversely, if extraction from any raw description of x produces descriptions of $N(x)$ and r_x with costs bounded by e_1 and e_2 , then

$$C_N(N(x)) \leq e_1(C_X(x)), \quad C_{\text{rel}}(x \mid N(x)) \leq e_2(C_X(x)).$$

Proof. The upper bound follows by assembling near-optimal descriptions of the normal datum and the residual datum. The two reverse inequalities follow by applying the extraction procedure to a near-optimal raw description of x . $\square$

Proposition 6.5 (Residual entropy). *Let $\mathcal{S} \subseteq \mathcal{X}$ be finite. For any normal map $N : \mathcal{S} \rightarrow N(\mathcal{S})$,*

$$\log |\mathcal{S}| \leq \log |N(\mathcal{S})| + \max_{y \in N(\mathcal{S})} \log |\text{Res}_N(y) \cap \mathcal{S}|.$$

Proof. The fibres of N partition $\mathcal{S}$, hence

$$|\mathcal{S}| \leq |N(\mathcal{S})| \cdot \max_y |\text{Res}_N(y) \cap \mathcal{S}|.$$

Taking logarithms gives the claim. $\square$

7 Separators and Lower Bounds

Normal forms are useful for lower bounds when their bounded images can be separated.

Let $\mathcal{N} : \Gamma_0 \rightsquigarrow \Gamma_1$ be a normal-form package. Define the budgeted normal image

$$S_{\mathcal{N},b} = \{n \in \mathcal{D}_1 : \exists d, \nu, \kappa_0(d) \leq b, \text{Comp}_{\mathcal{N}}(d, n, \nu)\}.$$

Definition 7.1 (Separator). *A separator for $S_{\mathcal{N},b}$ is a function*

$$s : \mathcal{D}_1 \rightarrow P$$

into a pointed preordered set $(P, 0)$ such that

$$s(n) = 0 \quad \text{for every } n \in S_{\mathcal{N},b}.$$

Theorem 7.2 (Separator lower bound). *Let s be a separator for $S_{\mathcal{N},b}$. If an object x has normal datum n_x and*

$$s(n_x) > 0,$$

then

$$C_{\Gamma_0}(x) > b.$$

Proof. If $C_{\Gamma_0}(x) \leq b$, then some bounded raw description of x compiles to an element of $S_{\mathcal{N},b}$ representing the same normal datum as n_x . Since s vanishes on the budgeted image, this contradicts $s(n_x) > 0$. $\square$

Definition 7.3 (Separation cost). *The separation cost of proving the lower bound for x at budget b through $\mathcal{N}$ is*

$$C_{\text{sep}}^{\mathcal{N}}(x, b) = \inf\{\kappa(s) + |\eta| : s(n_x) > 0, \eta \text{ verifies } s|_{S_{\mathcal{N},b}} = 0\}.$$

Thus a normal form may be conceptually perfect and still methodologically weak if the relevant bounded normal image is hard to separate. This is a genuine boundary of the method: algebraic non-membership, convex separation, integer programming, rigidity, and proof complexity reappear as domain-specific separation problems.

8 Fragmented and Effective Normal Forms

Many normal forms are complete only as an increasing tower. The following elementary theorem records the effective obstruction.

Let $P \subseteq \Sigma^*$ be a predicate with decidable fragments

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

such that

$$P = \bigcup_{r \geq 0} P_{\leq r}.$$

For $x \in P$, set

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

For finite balls, define

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

Theorem 8.1 (Fragmented degree theorem). *Assume finite alphabets, decidable fragments $P_{\leq r}$, and finite balls $\{|x| \leq n\}$. Then*

$$W_P \equiv_T P.$$

Consequently, if P is undecidable, then W_P has no computable total majorant.

Proof. First, $P \leq_T W_P$. Given x with $|x| = n$, compute $W_P(n)$ using the oracle, then decide whether

$$x \in P_{\leq W_P(n)}.$$

This works because $W_P(n)$ dominates the minimal fragment level of every positive instance of length at most n .

Conversely, $W_P \leq_T P$. Given n , enumerate all words of length at most n . Use the oracle for P to identify the positive instances. For each positive instance x , search over $r = 0, 1, 2, \dots$ until $x \in P_{\leq r}$. This terminates by assumption. Taking the maximum gives $W_P(n)$.

If W_P had a computable total majorant g , then membership in P would be decidable by testing $x \in P_{\leq g(|x|)}$, contradicting undecidability. $\square$

This theorem applies to fragmented normal-form towers, bounded proof search, bounded saturation, bounded finite quotient search, and bounded observable towers. Its content is not that the tower is useless, but that completeness may require a non-computable truncation function.

9 Minimal and Canonical Selectors

Part I attaches two decision problems to a filtered effective presentation system

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

The bounded-presentation problem P_Γ asks whether the E -class of d meets the finite cost ball $D_{\leq b}$. The equivalence problem EQ_Γ asks whether two descriptions lie in the same E -class.

Definition 9.1 (Minimal selector). *A minimal selector for Γ is a map*

$$N : D \rightarrow D$$

such that, for every $d \in D$,

$$N(d) E d, \quad \kappa(N(d)) = C_\Gamma(d).$$

It chooses a cheapest representative of the represented object, but it need not choose the same representative on equivalent inputs.

Definition 9.2 (Canonical minimal selector). *A canonical minimal selector is a minimal selector $N : D \rightarrow D$ such that, for all $d, e \in D$,*

$$d E e \iff N(d) = N(e).$$

It is a complete normal form with minimal presentation cost.

Theorem 9.3 (Selector criterion). *For a filtered effective presentation system:*

- (i) *a computable minimal selector exists if and only if P_Γ is decidable;*
- (ii) *a computable canonical minimal selector exists if and only if P_Γ and EQ_Γ are both decidable.*

Proof. Assume first that P_Γ is decidable. By Part I, C_Γ is computable. Given d , compute $c = C_\Gamma(d)$, enumerate the finite set $D_{\leq c}$, and search through stages $t = 0, 1, 2, \dots$ until some $e \in D_{\leq c}$ satisfies $d E_t e$. Such an e exists by definition of c . The first such e gives a computable minimal selector.

Conversely, if a computable minimal selector N exists, then

$$C_\Gamma(d) = \kappa(N(d))$$

is computable, so P_Γ is decidable.

For the canonical statement, assume P_Γ and EQ_Γ are decidable. Given d , compute $c = C_\Gamma(d)$, enumerate $D_{\leq c}$, and choose the lexicographically least $e \in D_{\leq c}$ such that $e E d$. This is computable because EQ_Γ is decidable. Equivalent inputs have the same finite set of cheapest representatives, hence the same chosen representative; non-equivalent inputs cannot choose the same representative. Thus the selector is canonical.

Conversely, if N is a computable canonical minimal selector, then $C_\Gamma(d) = \kappa(N(d))$ is computable, so P_Γ is decidable. Moreover,

$$d E e \iff N(d) = N(e),$$

and equality of finite strings is decidable. Hence EQ_Γ is decidable. $\square$

Remark 9.4 (Why this belongs to normal-form theory). *A cost-controlled compiler can produce useful representatives without resolving the whole equivalence relation. A minimal selector solves the optimization part of normalization. A canonical selector solves both optimization and classification. The theorem says exactly which effective decision problems are hidden in those two forms of normal access.*

10 Heighted and Degenerating Normal Forms

Normal-form compilation often interacts with arithmetic height. A normal form may be formally available, yet any representative realizing it may require large gauge data, large coefficients, or a high-complexity degeneration path. The following abstract package isolates the reusable mechanism.

Definition 10.1 (Heighted normal-form package). *A heighted normal-form package consists of projective varieties*

$$G, \quad N, \quad X$$

over $\overline{\mathbb{Q}}$, height functions

$$h_G, \quad h_N, \quad h_X$$

coming from fixed projective embeddings, and a rational action or realization map

$$a : G \times N \dashrightarrow X.$$

The space G is the gauge space, N is the normal-data space, and X is the object space. On each declared chart of the domain of a , the coordinate functions of a are represented by bihomogeneous forms of bidegree

$$(\delta_G, \delta_N).$$

A representation of $x \in X$ is a pair (g, n) in the declared domain with

$$x = a(g, n).$$

Its normal-form cost is

$$\text{NF}_{\mathcal{H}}(x; g, n) = \delta_G h_G(g) + \delta_N h_N(n),$$

and the minimal cost is the infimum over all such representations.

Theorem 10.2 (No free heighted normal forms). *For every heighted normal-form package $\mathcal{H}$, there is a constant $C_{\mathcal{H}}$, depending only on the chosen charts and coordinate forms, such that*

$$h_X(a(g, n)) \leq \delta_G h_G(g) + \delta_N h_N(n) + C_{\mathcal{H}}$$

for every declared representation (g, n) . Consequently,

$$\delta_G h_G(g) \geq h_X(x) - \delta_N h_N(n) - C_{\mathcal{H}}$$

whenever $x = a(g, n)$.

Proof. Choose projective coordinates on G, N, X . On a declared chart write

$$a(g, n) = [F_0(g, n) : \cdots : F_M(g, n)],$$

where the F_i are bihomogeneous of bidegree (δ_G, δ_N) and do not vanish simultaneously on the domain. The standard multihomogeneous height estimate gives

$$h_X(a(g, n)) \leq \delta_G h_G(g) + \delta_N h_N(n) + h(F) + C,$$

where F denotes the fixed tuple of coordinate forms. Absorb the fixed terms into $C_{\mathcal{H}}$, and rearrange. $\square$

Corollary 10.3 (Bounded normal data force high gauge). *If $x = a(g, n)$ and $h_N(n) \leq B_N$, then*

$$h_G(g) \geq \frac{h_X(x) - \delta_N B_N - C_{\mathcal{H}}}{\delta_G}.$$

Thus objects of unbounded height cannot be produced using both bounded-height normal data and bounded-height gauges.

Definition 10.4 (Degenerating normal-form representation). *Fix a projective compactification $\overline{G}$. A degenerating normal-form representation of $x \in X$ is a pair (γ, n) , where*

$$\gamma : \mathbb{P}^1 \rightarrow \overline{G}$$

is a homogeneous algebraic path of degree D and

$$x = \lim_{t \rightarrow 0} a(\gamma(t), n)$$

exists in X . Its degenerating cost has the form

$$\text{DNF}_{\mathcal{H}}(x; \gamma, n) = \delta_G h(\gamma) + \delta_N h_N(n) + C \log(D + 2),$$

with C depending only on the package.

Theorem 10.5 (No free degenerating normal forms). *For every heighted normal-form package there is a constant $C_{\mathcal{H}}$ such that every degenerating representation (γ, n) of x satisfies*

$$h_X(x) \leq \delta_G h(\gamma) + \delta_N h_N(n) + C_{\mathcal{H}} \log(D + 2) + C_{\mathcal{H}}.$$

Equivalently,

$$\delta_G h(\gamma) + C_{\mathcal{H}} \log(D + 2) \geq h_X(x) - \delta_N h_N(n) - C_{\mathcal{H}}.$$

Proof. Write the path as

$$\gamma = [\Gamma_0(S, T) : \cdots : \Gamma_r(S, T)]$$

with homogeneous coordinate forms of common degree D . Substitute these forms into the bihomogeneous coordinate tuple defining a . The coefficients of the resulting homogeneous forms are sums of products involving δ_G coefficients from γ , δ_N coordinates from n , and fixed coefficients of the map a . The number of summands is polynomial in D , so the logarithmic height of the substituted tuple is bounded by

$$\delta_G h(\gamma) + \delta_N h_N(n) + O(\log(D + 2)) + O(1).$$

Passing to the limit at $t = 0$ amounts to taking the first nonzero coefficient vector after removing a common vanishing order. This operation does not increase the height beyond the coefficient-height bound up to a fixed additive constant. The displayed estimate follows. $\square$

Definition 10.6 (Height-controlled section). *Let*

$$q : X^{ss} \rightarrow \mathcal{M}$$

be a quotient, moduli, or orbit-class map. A rational section

$$s : U \dashrightarrow X^{ss}$$

over $U \subseteq \mathcal{M}$ is height-controlled if there are constants A_s, B_s such that

$$h_X(s(m)) \leq A_s h_{\mathcal{M}}(m) + B_s$$

for all m in the declared domain.

Theorem 10.7 (Representative-moduli gap). *Let $a : G \times X \dashrightarrow X$ be a heightened gauge action of bidegree (δ_G, δ_X) , and let s be a height-controlled section of $q : X^{ss} \rightarrow \mathcal{M}$. If*

$$x = a(g, s(q(x))),$$

then

$$h_X(x) \leq \delta_G h_G(g) + \delta_X A_s h_{\mathcal{M}}(q(x)) + O(1).$$

Consequently,

$$h_G(g) \geq \frac{h_X(x) - \delta_X A_s h_{\mathcal{M}}(q(x)) - O(1)}{\delta_G}.$$

Proof. Apply the no-free heightened normal-form inequality to

$$x = a(g, s(q(x))).$$

The height-controlled section bounds $h_X(s(q(x)))$ by $A_s h_{\mathcal{M}}(q(x)) + B_s$. Substitute and rearrange. $\square$

Proposition 10.8 (Controlled transfer of heightened normal forms). *Let $\mathcal{H}_P$ and $\mathcal{H}_Q$ be heightened normal-form packages. Suppose rational maps*

$$T_G : G_P \dashrightarrow G_Q, \quad T_N : N_P \dashrightarrow N_Q, \quad T_X : X_P \dashrightarrow X_Q$$

intertwine the realization maps:

$$T_X(a_P(g, n)) = a_Q(T_G(g), T_N(n)),$$

and suppose T_G, T_N control heights linearly up to additive constants. Then the minimal heightened normal-form cost in $\mathcal{H}_Q$ is bounded above by an affine function of the minimal heightened normal-form cost in $\mathcal{H}_P$.

Proof. Choose a representation $x = a_P(g, n)$. The intertwining identity gives a representation of $T_X(x)$ in $\mathcal{H}_Q$. The height bounds for T_G and T_N control its Q -cost by an affine function of the P -cost of (g, n) . Taking the infimum over all P -representations of x gives the claim. $\square$

11 Rewriting and Saturation

Rewriting systems are a standard source of normal forms. The presentation-theoretic point is to keep the cost growth visible.

Definition 11.1 (Costed rewriting system). *A costed abstract rewriting system consists of data*

$$(\mathcal{D}, \rightarrow, \rho, \kappa, h, \sigma)$$

such that

$$d \rightarrow e \implies \rho(d) = \rho(e),$$

$$h(e) < h(d),$$

and

$$\kappa(e) \leq \sigma(\kappa(d)).$$

Here h is a termination height and σ controls cost growth.

Theorem 11.2 (Budgeted Newman principle). *Suppose the rewrite system is terminating by h , locally confluent, and $h(d) \leq H(\kappa(d))$. Then each d has a unique normal form in its rewrite component. If $\kappa(d) \leq b$, the normal form has cost at most*

$$\sigma^{H(b)}(b).$$

If each rewrite step has verification data of cost at most $c(b)$, the full verification data have cost at most

$$H(b) \cdot c(\sigma^{H(b)}(b)).$$

Proof. Termination and local confluence imply confluence by Newman's lemma, so each rewrite component has a unique normal form. Any reduction path has length at most $H(b)$. Iterating the one-step cost bound gives $\sigma^{H(b)}(b)$. Concatenating the stepwise verification data gives the displayed verification bound. $\square$

12 The Normal-Form Landscape

It is useful to organize normal-form packages for a fixed presentation system without pretending that this organization is automatically a rigid geometric object.

Definition 12.1 (Normal-form landscape). *For a presentation system Γ , the normal-form landscape*

$$\text{NFLand}(\Gamma)$$

is the collection of normal-form packages out of Γ , considered up to mutual simulation with declared overhead.

There is a natural domination relation:

$$\mathcal{M} \succeq \mathcal{N}$$

if the normal data of $\mathcal{N}$ can be computed from the normal data of $\mathcal{M}$ with controlled overhead. In this case $\mathcal{M}$ is at least as informative as $\mathcal{N}$.

The landscape contains overlapping regions:

representative, observable, relative, stratified, stacky, approximate, fragmented.

One may also study Pareto-optimal normal forms with respect to informativeness, cost, verification, residual size, and separation cost.

Remark 12.2. *The purpose of the landscape language is bookkeeping. It is a way to compare normalizers, not a claim that every presentation system comes with a canonical spectrum. In examples with additional topology, category theory, or sheaf structure, one can enrich this bookkeeping into a more geometric object.*

13 Local Normal Forms and Descent Obstructions

Many normal forms exist only locally. Eigenvector choices, gauge choices, coordinate gauges, local trivializations, and normal forms by strata all have this flavour.

Assume a space or groupoid of objects $\mathcal{X}$ is covered by regions

$$\mathcal{X} = \bigcup_i U_i$$

and that on each U_i there is a local normal-form package

$$\mathcal{N}_i : U_i \rightsquigarrow T_i.$$

On overlaps $U_{ij} = U_i \cap U_j$, transition data compare T_i and T_j . On triple overlaps, the failure of the transitions to compose strictly lies in an automorphism sheaf or groupoid of local normal data.

Proposition 13.1 (Descent obstruction, schematic form). *In a setting where the local normal data form a sheaf or stack with transition data g_{ij} , a strict global normal form compatible with the local ones exists only if the corresponding descent obstruction vanishes. In the abelian sheaf case this obstruction is represented by a Čech cocycle; in the nonabelian case it is represented by a gerbe-type class.*

Proof. This is the standard descent criterion for gluing local objects with transition functions. If the cocycle is a coboundary, local gauge corrections make the transitions compatible and the local data glue. If it is not a coboundary, no strict global object with those transition data exists. $\square$

This section is included to locate normal forms with unavoidable choices. The spectral application below is the simplest case: local eigenvector choices exist, but global continuous sign choices do not.

14 Spectral Canonicalization

We now give a concrete application. It concerns spectral graph methods, eigenvector positional encodings, and the instability of global eigenbasis choices.

Let

$$\text{Sym}_N$$

be the space of real symmetric $N \times N$ matrices. Fix real numbers

$$\lambda_1 < \lambda_2 < \cdots < \lambda_N$$

with

$$\lambda_{i+1} - \lambda_i \geq \delta > 0.$$

Define the fixed-spectrum orbit

$$\mathcal{O}_\lambda = \{Q \text{diag}(\lambda_1, \dots, \lambda_N) Q^T : Q \in O(N)\}.$$

Every matrix in $\mathcal{O}_\lambda$ has simple spectrum and eigengap at least δ .

Definition 14.1 (Eigenbasis canonicalizer). *An eigenbasis canonicalizer is a map*

$$C : \mathcal{O}_\lambda \rightarrow O(N)$$

such that

$$A C(A) = C(A) \text{diag}(\lambda_1, \dots, \lambda_N).$$

The columns of $C(A)$ are ordered unit eigenvectors. Since every eigenspace is a line, such a map is precisely a global sign choice for the ordered eigenline decomposition.

Theorem 14.2 (Fixed-gap spectral canonicalization instability). *For $N \geq 2$, no eigenbasis canonicalizer*

$$C : \mathcal{O}_\lambda \rightarrow O(N)$$

can be continuous.

More strongly, for any deterministic eigenbasis canonicalizer C , even if discontinuous, there exist matrices

$$A_m \rightarrow A$$

inside $\mathcal{O}_\lambda$, all with eigengap at least δ , such that for some eigenvector index i ,

$$\|u_i(A_m) - u_i(A)\| \rightarrow 2,$$

where $u_i(A)$ is the i -th column of $C(A)$. Consequently no global eigenbasis canonicalizer can be uniformly continuous, Holder continuous, or Lipschitz on the fixed-gap simple-spectrum orbit.

Proof. It is enough to restrict to a two-dimensional rotation. Let

$$R_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}.$$

For $\theta \in [0, \pi]$, define

$$A_\theta = R_\theta \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} R_\theta^T \oplus \text{diag}(\lambda_3, \dots, \lambda_N).$$

Then $A_0 = A_\pi$, because $R_\pi = -I_2$. All matrices A_θ lie in $\mathcal{O}_\lambda$, and all have the same eigengap.

The first eigenline of A_θ is $[R_\theta e_1]$. As θ runs from 0 to π , this line gives the standard nontrivial loop in $\mathbb{R}P^1$. A continuous choice of unit eigenvector would lift this loop through the antipodal covering

$$S^1 \rightarrow \mathbb{R}P^1.$$

Starting at e_1 , the lifted path ends at $-e_1$, not at e_1 . But $A_0 = A_\pi$, so a single-valued continuous canonicalizer would require the endpoint to equal the starting value. This contradiction proves non-existence of a continuous C .

For the stronger statement, restrict any deterministic canonicalizer to the first eigenline along the same projective loop. It gives a section of the double cover $S^1 \rightarrow \mathbb{R}P^1$. Since no continuous section exists, the section is discontinuous somewhere. At a discontinuity, choose a convergent sequence of eigenlines $\ell_m \rightarrow \ell$. Passing to a subsequence, the chosen unit vectors converge to the opposite sign of the chosen vector over ℓ . Hence the norm difference tends to 2. Realizing these eigenlines by matrices in $\mathcal{O}_\lambda$ gives the required sequence. $\square$

14.1 Stable replacement

The complete eigenbasis datum

$$A \mapsto (\lambda_i, u_i)_i$$

is a set-theoretic normal form on the simple-spectrum orbit, but it is not stable. The stable normal form is instead

$$A \mapsto (\lambda_i, P_i)_i, \quad P_i = u_i u_i^T.$$

The residual fibre is the sign torsor $\{\pm 1\}^N$. If eigenvalues have multiplicities m_λ , the stable datum is

$$A \mapsto \{(\lambda, P_\lambda)\}_\lambda,$$

and the residual gauge group is

$$\prod_\lambda O(m_\lambda).$$

This is a stacky normal form.

Theorem 14.3 (Universal projector factorization). *Let $\text{St}(m, N)$ be the Stiefel manifold and $\text{Gr}(m, N)$ the Grassmannian. Let*

$$\pi : \text{St}(m, N) \rightarrow \text{Gr}(m, N), \quad U \mapsto UU^T$$

be the quotient map. If Y is a topological space and

$$F : \text{St}(m, N) \rightarrow Y$$

is continuous and right $O(m)$ -invariant, then there is a unique continuous map

$$\tilde{F} : \text{Gr}(m, N) \rightarrow Y$$

such that

$$F = \tilde{F} \circ \pi.$$

Proof. The map π is the quotient map for the free right action of $O(m)$ on $\text{St}(m, N)$. Since F is constant on $O(m)$ -orbits, it factors uniquely through the quotient. Continuity of the induced map follows from the universal property of quotient maps. $\square$

Thus projector-valued spectral normal forms are not just a convenient workaround. They are universal among continuous basis-invariant encodings. In graph-learning language, sign- and basis-invariant spectral methods align with the stable normal form, while complete eigenbasis canonicalization crosses a topological obstruction to global stability.

15 Further Templates

This section records several patterns already used elsewhere in this collection. They are not presented as new applications here; their role is to show how the normal-form language organizes existing arguments.

15.1 Sparse factorization

A local linear data structure for a workload matrix $A \in \mathbb{F}^{m \times n}$, with storage s , query locality q , and update locality u , can be normalized to a factorization

$$A = RM,$$

where $R \in \mathbb{F}^{m \times s}$, $M \in \mathbb{F}^{s \times n}$, each row of R has support at most q , and each column of M has support at most u . The budgeted normal image is the set of such sparse factorizations. Any separator proving that A is outside that image gives a lower bound for the original data-structure presentation.

15.2 Boolean lifts

For a Boolean function $f : \{0, 1\}^n \rightarrow \{-1, 1\}$, a polynomial lift model may impose a coefficient alphabet, a degree bound, height bounds, sparsity bounds, or homogeneity. The finite-difference or Mobius coefficient vector is a normal datum. The lift model imposes lattice and capacity constraints on this datum. Lower bounds become lattice-capacity separation problems.

15.3 Algebraic proof degree

A static algebraic proof identity

$$h = \sum_i g_i f_i$$

with degree bound $\deg(g_i f_i) \leq D$ can be normalized to homogeneous membership

$$x_0^{D-\deg h} h^h \in J.$$

The obstruction is torsion in

$$(J : x_0^\infty)/J.$$

Thus proof-degree lower bounds become separation problems in a filtered homogeneous module.

15.4 Class-two p -groups

For special class-two exponent- p groups, the commutator is an alternating bimap

$$\beta : \Lambda^2 V \rightarrow W.$$

Words admit a linear-quadratic normal datum, and ordinary word-measure observables factor through rank-profile data such as

$$\lambda \mapsto \text{rank}(\lambda \circ \beta).$$

Large residual fibres of this observable quotient produce non-isomorphic groups indistinguishable by the chosen observables.

16 Endpoint

The framework reaches a real boundary: a normal-form package can expose a budgeted image, and it can identify the normal datum whose membership must be tested, but it cannot automatically produce the optimal separator in every domain. The hard mathematics reappears in the separation problem.

The practical value is therefore diagnostic. A normal-form analysis tells us:

what was rigidified,
what remains in the residual fibre,
what the verification data control,
which observables are preserved,
and what must be separated to prove a lower bound.

17 Conclusion

This note replaces the question

“what is the normal form?”

with the more structured question

“what normal forms exist, what do they preserve, what do they lose,
what do they cost, and what can they separate?”

The answer is a normal-form calculus:

normal datum + verification data + residual fibre + separator + descent obstruction.

The spectral application shows the calculus at work. Complete eigenbasis data exist set-theoretically on the simple-spectrum locus, but no global stable eigenbasis canonicalizer exists even under a fixed eigengap. The stable normal form is projector-valued and stacky. Normalization does not remove complexity; it redistributes it among rigid data, residual fibres, choices, verification data, and separation cost.

References

- [1] L. Blanchi, *Presentation Theory I: Description Systems, Costs, and Observable Fibres*, Presentation Theory note, 2026.
- [2] L. Blanchi, *Presentation Theory II: Controlled Transfer, Observable Budgets, and the Geometry of Fibres*, Presentation Theory note, 2026.
- [3] L. Blanchi, *Atlas of Controlled Transfer Packages*, Presentation Theory note, 2026.
- [4] M. H. A. Newman, *On theories with a combinatorial definition of equivalence*, Annals of Mathematics 43 (1942), 223–243.
- [5] D. Lim, J. Robinson, L. Zhao, T. Smidt, S. Sra, H. Maron, and S. Jegelka, *Sign and Basis Invariant Networks for Spectral Graph Representation Learning*, arXiv:2202.13013, 2022.
- [6] S. Hordan, N. Dym, and T. Seppelt, *Weisfeiler-Leman Is Incomplete on Simple Spectrum Graphs, so Canonicalize Them*, arXiv:2605.23446, 2026.
- [7] J. Ma, Y. Wang, and Y. Wang, *Laplacian Canonization: A Minimalist Approach to Sign and Basis Invariant Spectral Embedding*, arXiv:2310.18716, 2023.