Invisible Flats and the Cost of Marked Normal Forms

Luca Blanchi

June 2026

Abstract

This note proves a coarse rate-distortion obstruction for reconstructing full markings from finite-depth Torelli, Johnson, homological, level, or similar partial data. The general mechanism is simple and robust. If an observable is constant on a quasi-isometrically embedded abelian orbit $\mathbb{Z}^r \cdot x$, then every inverse normal form from the observable back to a marked presentation has worst-case and average error $\Omega(rM)$ on boxes of side length M. The same invisible flat contains $(M/R)^{\Omega(r)}$ many R-separated presentations, so reconstruction at resolution R requires $\Omega(r \log(M/R))$ auxiliary bits.

For mapping class groups this abstract obstruction is realized by standard multitwist and bounded-support Johnson-filtration flats. Unmarked curve data, unmarked Jacobian data, and finite marking rigidifications leave invisible flats of rank $3g-3$ on the generic closed genus g locus. Marked homological data leave the Torelli kernel, which contains the separating multitwist flat of rank $2g-3$, and the same rank persists after adding the first Johnson quotient. For every fixed Johnson depth k on $S_{g,1}$, the residual kernel contains undistorted abelian subgroups of rank at least $c_k g - O_k(1)$.

Thus finite-depth Torelli or Johnson data may rigidify finite stabilizers or determine an unmarked object, but they do not give low-cost marked normal forms. The result is a presentation-theoretic lower bound: the large-scale geometry of the fibre forces an arithmetic of missing marking information.

1 Introduction

A moduli problem often separates an object from a presentation of that object. For a smooth curve C of genus g, one may remember only the unmarked curve, or one may remember a marking $f : S_g \to C$. Torelli-type theorems can determine the unmarked object very strongly. They do not automatically produce a low-cost marked presentation.

This note quantifies the obstruction. The relevant cost is not the difficulty of recognizing the curve. It is the coarse cost of choosing a full marking from data that are constant on a large residual gauge group.

The guiding principle is:

large invisible flats in a fibre force large normal-form cost.

If an observable is constant on an undistorted orbit $\mathbb{Z}^r \cdot x$, then an inverse normal form must choose one output for a box of mutually far marked presentations. On an r-dimensional box of side length M, this forces error of order rM, and it forces $r \log(M/R)$ bits of extra data to reconstruct the marking at resolution R.

The mapping class group provides natural high-rank invisible flats. Pants multitwists give rank $3g-3$ flats in the full marking fibre. Separating multitwists give rank $2g-3$ flats in the Torelli and Johnson-kernel fibres. For deeper fixed Johnson levels, bounded-support local elements can be placed on linearly many disjoint subsurfaces, giving invisible flats of rank $\Omega_k(g)$.

Theorem 1.1 (Informal main theorem). *Finite level, Torelli, first-Johnson, and fixed-depth Johnson observables do not admit low-cost inverse normal forms to full markings. More precisely, their fibres contain undistorted abelian flats of rank linear in genus, and therefore every inverse normal form has linear-genus coarse error and linear-genus missing bit complexity on large boxes.*

The theorem is compatible with classical Torelli. Torelli determines the unmarked curve from its principally polarized Jacobian on the smooth locus. The present result concerns the additional cost of recovering a marking.

2 Coarse Marked Reconstruction

Let X be a metric space of marked presentations, let Γ be a finitely generated group acting on X by isometries or uniformly quasi-isometric maps, and let

$$q : X \to Z$$

be an observable. The fibre $q^{-1}(q(x))$ is the set of presentations indistinguishable from x by the observable.

Definition 2.1 (Inverse normal form). *An inverse normal form for q is an arbitrary map*

$$N : q(X) \to X.$$

No continuity, measurability, definability, or computability hypothesis is imposed. The lower bounds below therefore apply to all selectors.

Definition 2.2 (Invisible abelian quasi-flat). *Let $A \cong \mathbb{Z}^r$ be generated by $a_1, \dots, a_r$. The orbit $A \cdot x$ is a q-invisible abelian quasi-flat of rank r if*

$$q(ax) = q(x) \qquad \text{for all } a \in A$$

and the orbit map

$$\mathbb{Z}^r \to X, \qquad n = (n_1, \dots, n_r) \mapsto a_1^{n_1} \cdots a_r^{n_r} x$$

is a quasi-isometric embedding. Equivalently, there are constants $\alpha > 0$, $\beta \geq 0$ such that

$$d_X(a^n x, a^m x) \geq \alpha |n - m|_1 - \beta$$

for all $n, m \in \mathbb{Z}^r$.

Definition 2.3 (Invisible flat rank). *The invisible flat rank of q at x is*

$$\operatorname{ifr}_x(q) = \sup\{r : \ q^{-1}(q(x)) \text{ contains a } q\text{-invisible abelian quasi-flat of rank } r\}.$$

The fibre flat rank $\operatorname{ffr}_x(q)$ is defined similarly, but allows arbitrary quasi-isometrically embedded copies of $\mathbb{Z}^r$ in the fibre rather than only abelian group orbits. Always $\operatorname{ffr}_x(q) \geq \operatorname{ifr}_x(q)$.

Proposition 2.4 (Coarse invariance of fibre flat rank). *Let $q : X \to Z$ and $q' : X' \to Z'$ be observables. Suppose that $F : X \to X'$ is a quasi-isometry and that*

$$q(x) = q(y) \quad \implies \quad q'(F(x)) = q'(F(y))$$

on the region under consideration. If $q^{-1}(q(x))$ contains a quasi-isometrically embedded copy of $\mathbb{Z}^r$, then $q'^{-1}(q'(F(x)))$ also contains such a copy.

Proof. Compose the quasi-isometric embedding $\mathbb{Z}^r \to q^{-1}(q(x))$ with F. The result remains a quasi-isometric embedding, and the fibre-respecting hypothesis puts its image inside one q'-fibre. □

3 Invisible Flats and Rate-Distortion

Let $A \cong \mathbb{Z}^r$ be generated by $a_1, \ldots, a_r$. For

$$n = (n_1, \ldots, n_r)$$

write $a^n = a_1^{n_1} \cdots a_r^{n_r}$, and put

$$Q_M = [-M, M]^r \cap \mathbb{Z}^r.$$

Theorem 3.1 (Invisible flat rate-distortion)**.** *Let $A \cdot x$ be a q-invisible abelian quasi-flat, and suppose*

$$d_X(a^n x, a^m x) \geq \alpha |n - m|_1 - \beta.$$

Then every inverse normal form $N : q(X) \to X$ satisfies

$$\sup_{n \in Q_M} d_X(N(q(x)), a^n x) \geq \alpha r M - \frac{\beta}{2}.$$

Moreover, if n is uniformly distributed in Q_M, then

$$\mathbb{E}_{n \in Q_M} d_X(N(q(x)), a^n x) \geq \alpha r \frac{M(M+1)}{2M+1} - \frac{\beta}{2}.$$

In particular, for $M \geq 1$, the average error is

$$\Omega_{\alpha,\beta}(rM).$$

Proof. Let $n_+ = (M, \ldots, M)$ and $n_- = (-M, \ldots, -M)$. Since $A \cdot x$ is invisible,

$$q(a^{n_+} x) = q(a^{n_-} x) = q(x).$$

The inverse normal form gives the same output for both points. By the triangle inequality, at least one of the two distances from $N(q(x))$ is at least half of $d_X(a^{n_+} x, a^{n_-} x)$. The quasi-isometric lower bound gives

$$d_X(a^{n_+} x, a^{n_-} x) \geq 2\alpha r M - \beta.$$

This proves the worst-case estimate.

For the average estimate, pair n with $-n$. Again the two points have the same observable value, so

$$d_X(N(q(x)), a^n x) + d_X(N(q(x)), a^{-n} x) \geq d_X(a^n x, a^{-n} x) \geq 2\alpha |n|_1 - \beta.$$

Averaging over Q_M gives

$$\mathbb{E} d_X(N(q(x)), a^n x) \geq \alpha\, \mathbb{E}|n|_1 - \frac{\beta}{2}.$$

Since

$$\mathbb{E}_{j \in [-M,M] \cap \mathbb{Z}} |j| = \frac{M(M+1)}{2M+1},$$

we have

$$\mathbb{E}_{n \in Q_M} |n|_1 = r \frac{M(M+1)}{2M+1}.$$

The claim follows. □

Corollary 3.2 (Packing and missing bits). *Under the hypotheses of the theorem, the set*

$$F_M = \{a^n x : n \in Q_M\}$$

contains an R-separated subset of cardinality at least

$$C_\alpha \left(\frac{M}{R+\beta+1}\right)^r$$

for $M \gg R+\beta$. Consequently, any reconstruction procedure which is given $q(x)$ and an auxiliary message of b bits, and which must reconstruct every point of F_M up to error R, requires

$$b \geq r \log_2\left(\frac{M}{R+\beta+1}\right) - O_\alpha(r).$$

Proof. Choose a coordinate subgrid of Q_M with spacing

$$s = \left\lceil \frac{2R+\beta+1}{\alpha} \right\rceil.$$

Distinct points in this subgrid have ℓ^1-distance at least s, hence their orbit points are R-separated. The subgrid has cardinality

$$\gtrsim_\alpha \left(\frac{M}{R+\beta+1}\right)^r.$$

For the bit bound, a b-bit message gives at most 2^b possible outputs over the fixed observable value. These outputs must form an R-net for F_M, so 2^b is at least the packing number at a comparable radius. Taking logarithms proves the estimate. □

Remark 3.3 (Flat-local sharpness). *On the invisible flat itself the bit estimate is sharp up to $O(r)$: cover the coordinate box Q_M by ℓ^1-balls of radius comparable to R and transmit the index of the ball. This gives an upper bound*

$$r \log(M/R) + O(r)$$

for reconstructing points on the flat. The full fibre may have richer nonabelian geometry, so the flat-local upper bound need not describe the entire reconstruction problem.

4 Mapping Class Group Input

Let S_g be a closed oriented surface of genus g, and let $S_{g,1}$ be a genus g surface with one boundary component fixed pointwise. Write

$$\Gamma_g = \mathrm{Mod}(S_g), \qquad \Gamma_{g,1} = \mathrm{Mod}(S_{g,1}).$$

Let $\mathcal{M}(S)$ denote the marking graph. The action of $\mathrm{Mod}(S)$ on $\mathcal{M}(S)$ is geometric, so $\mathcal{M}(S)$ is quasi-isometric to the mapping class group with any word metric.

Proposition 4.1 (Multitwist flats). *If $\gamma_1, \ldots, \gamma_r$ are pairwise disjoint, pairwise non-isotopic essential simple closed curves, then the twists*

$$T_{\gamma_1}, \ldots, T_{\gamma_r}$$

generate a subgroup $\mathbb{Z}^r$, and its orbit in the marking graph is quasi-isometrically embedded.

Proof. The twists commute and are independent because the curves are disjoint and non-isotopic. The Masur-Minsky distance formula detects the exponent of T_{γ_i} in the annular projection to the annulus around γ_i. The annular contributions for disjoint curves add, giving a lower bound

$$d_{\mathrm{mark}}\Big(x, T_{\gamma_1}^{n_1} \cdots T_{\gamma_r}^{n_r} x\Big) \geq c \sum_i |n_i| - C.$$

The reverse inequality is immediate from applying the twists successively. □

We use the following standard rank inputs. The maximal rank of a free abelian subgroup of $\mathrm{Mod}(S_g)$ is $3g-3$, realized by a pants multitwist subgroup [BLM, FLM]. The Torelli group

$$\mathcal{I}_g = \ker(\mathrm{Mod}(S_g) \to \mathrm{Sp}_{2g}(\mathbb{Z}))$$

contains separating twists, and Vautaw's rank theorem gives maximal abelian rank $2g-3$ in $\mathcal{I}_g$, realized by separating multitwists [Vau]. The Johnson kernel $\mathcal{K}_g$, the kernel of the first Johnson homomorphism on $\mathcal{I}_g$, contains the subgroup generated by separating twists [Joh, Put].

For fixed Johnson depth we use a bounded-support locality input. Let $\mathcal{J}_{g,1}(k)$ be the k-th Johnson filtration term, with the convention $\mathcal{J}_{g,1}(1) = \mathcal{I}_{g,1}$. For every fixed k, the Church-Putman bounded-support theorem gives a constant G_k, independent of g, such that $\mathcal{J}_{g,1}(k)$ is generated by elements supported on homologically standard subsurfaces of genus at most G_k [CP]. Standard local nontriviality constructions, for instance the pseudo-Anosov elements in deep Johnson filtration terms constructed in [MP], give an infinite-order mapping class η_k supported on a surface of genus h_k, with $\eta_k \in \mathcal{J}(k)$, whose extensions by identity lie in $\mathcal{J}_{g,1}(k)$.

Lemma 4.2 (Disjoint local Johnson elements give flats)**.** *Fix k. Suppose $\eta_k \in \mathcal{J}(k)$ is an infinite-order element supported on a homologically standard subsurface W_k of genus h_k. If $W_{k,1}, \ldots, W_{k,r}$ are pairwise disjoint homologically standard copies of W_k in $S_{g,1}$, then, after replacing the corresponding extensions by positive powers if necessary, they generate an undistorted subgroup*

$$\mathbb{Z}^r \leq \mathcal{J}_{g,1}(k).$$

Proof. The supported mapping classes commute because their supports are disjoint. After passing to powers they are pure on their supports. A relation among the extensions restricts on each $W_{k,i}$ to a power of η_k, hence all exponents vanish because η_k has infinite order.

For undistortedness, each infinite-order pure mapping class has either a twisting component detected by an annular projection or a pseudo-Anosov component detected by a nonannular subsurface projection. These detecting subsurfaces lie in the disjoint supports, so the corresponding terms add in the Masur-Minsky distance formula. This gives the required linear lower bound in the sum of the exponents. □

5 Finite and Torelli-Level Observables

The first applications are closed-surface statements on the generic smooth locus. Finite automorphism groups are ignored in the usual coarse sense: they do not change the existence or rank of undistorted flats.

Theorem 5.1 (Unmarked curve and Jacobian data)**.** *For a generic closed genus $g \geq 3$ curve, the observable which remembers only the unmarked curve has invisible flat rank $3g-3$. The same is true for the observable which remembers only the unmarked principally polarized Jacobian. Therefore every inverse normal form from either observable to a full marking has worst-case and average error $\Omega(gM)$ on boxes of side length M, and requires $\Omega(g \log(M/R))$ auxiliary bits at resolution R.*

Proof. Forgetting the marking leaves the full marking fibre, coarsely modeled by $\mathrm{Mod}(S_g)$. A pants decomposition gives an invisible multitwist flat of rank $3g-3$, and no abelian subgroup of $\mathrm{Mod}(S_g)$ has larger rank. On the generic smooth locus, Torelli identifies the fibre of the unmarked Jacobian observable with the same unmarked curve fibre up to finite stabilizers. The rate-distortion bounds follow from the general theorem. □

Theorem 5.2 (No finite rigidification). *Let $q : \mathcal{M}(S_g) \to Z$ be an observable whose dependence on the marking factors through a finite quotient of $\mathrm{Mod}(S_g)$. Equivalently, suppose there is a finite-index subgroup $K \leq \mathrm{Mod}(S_g)$ such that*

$$q(kx) = q(x) \qquad \text{for all } k \in K.$$

Then

$$\mathrm{ifr}(q) = 3g - 3.$$

In particular, finite level structures may remove finite stabilizers, but they do not reduce the coarse cost of recovering a full marking.

Proof. Let $P \cong \mathbb{Z}^{3g-3}$ be a pants multitwist subgroup. Since K has finite index in $\mathrm{Mod}(S_g)$, the intersection $K \cap P$ has finite index in P, hence has the same rank. It is invisible for q and remains undistorted. The upper bound $3g-3$ is the maximal abelian rank of $\mathrm{Mod}(S_g)$. □

Corollary 5.3 (Homological level). *For every $m \geq 2$, the observable recording homological level m data leaves an invisible flat of rank $3g-3$. Indeed, the level subgroup has finite index, and powers T_γ^m of the twists in a pants decomposition act trivially on*

$$H_1(S_g; \mathbb{Z}/m\mathbb{Z}).$$

Theorem 5.4 (Marked homological data and the first Johnson quotient). *The observable whose residual gauge is the Torelli group has invisible flat rank*

$$2g - 3.$$

The same rank remains after adding the first Johnson quotient, so that the residual gauge is the Johnson kernel $\mathcal{K}_g$. Hence these data still force $\Omega(gM)$ coarse marked reconstruction error and $\Omega(g \log(M/R))$ auxiliary bits on large boxes.

Proof. The Torelli kernel contains the separating multitwist subgroup associated to a maximal family of disjoint separating curves. This gives an undistorted invisible flat of rank $2g-3$. Vautaw's theorem gives the matching upper bound for abelian subgroups of the Torelli group.

The Johnson kernel contains all separating twists, while it is contained in the Torelli group. Therefore the same lower and upper bounds apply after adding the first Johnson quotient. □

6 Fixed-Depth Johnson Data

We now work on $S_{g,1}$, where the Johnson filtration is most naturally defined. A depth-k Johnson observable is an observable

$$q_k : \mathcal{M}(S_{g,1}) \to Z_k$$

whose invisible gauge contains $\mathcal{J}_{g,1}(k)$.

Theorem 6.1 (Linear invisible flats in fixed Johnson depth). *For every fixed $k \geq 1$, there are constants $c_k > 0$ and $C_k \geq 0$ such that, for all sufficiently large g, the group $\mathcal{J}_{g,1}(k)$ contains an undistorted abelian subgroup*

$$A_{g,k} \cong \mathbb{Z}^{r_{g,k}}, \qquad r_{g,k} \geq c_k g - C_k.$$

Proof. Choose the local infinite-order element η_k supported on a homologically standard subsurface W_k of genus h_k described above. Embed pairwise disjoint copies of W_k into $S_{g,1}$. The number of such copies is at least

$$\left\lfloor \frac{g}{h_k} \right\rfloor - O_k(1).$$

Extend η_k by identity on each copy. Naturality of the Johnson filtration puts all these extensions in $\mathcal{J}_{g,1}(k)$. The preceding disjoint-support lemma shows that suitable positive powers generate an undistorted free abelian subgroup of the same rank. Taking

$$c_k = \frac{1}{2h_k}$$

and increasing C_k if needed gives the displayed linear lower bound. □

Corollary 6.2 (No fixed-depth marked reconstruction)**.** *Let q_k be a depth-k Johnson observable on $\mathcal{M}(S_{g,1})$. Every inverse normal form*

$$N : q_k(\mathcal{M}(S_{g,1})) \to \mathcal{M}(S_{g,1})$$

has worst-case and average error

$$\Omega_k(gM)$$

on boxes of side length M. Its invisible fibre has packing entropy at least

$$(M/R)^{\Omega_k(g)}$$

in the coarse regime $M \gg R$, and reconstruction at resolution R requires

$$\Omega_k(g \log(M/R))$$

auxiliary bits.

Proof. The invisible gauge of q_k contains $\mathcal{J}_{g,1}(k)$, hence contains the abelian quasi-flat from the theorem. Apply the invisible flat rate-distortion theorem and the packing corollary. □

7 Presentation-Theoretic Interpretation

The statements above are naturally presentation-theoretic. The marked object is a presentation, the mapping class group is a gauge group of presentation changes, and the observable records partial data. A normal form is a section of the observable map. Invisible flat rank measures how much coarse presentation geometry remains inside a fibre.

This is the same principle as in no-free-degeneration results, but in a purely coarse-geometric setting. A normal form does not become low-cost merely because the unmarked object is determined. The missing marking information is measured by large boxes inside the fibre.

The fixed-depth Johnson theorem is the main structural application. It says that making the observable deeper in a fixed Johnson tower does not remove the linear-genus Euclidean sector of the marking fibre. To make marked reconstruction low-cost, the depth or the auxiliary marking information must grow with the genus or with the desired scale.

8 Boundary Conventions and Questions

The sharp low-depth closed-surface statements use standard closed-surface rank theorems. The fixed-depth Johnson statement is written for $S_{g,1}$ because the Johnson filtration and the bounded-support theorem have their cleanest form there. Closed-surface variants follow whenever the observable's residual gauge contains the images of the local subgroups constructed above under the relevant capping or forgetting map.

Several refinements remain natural.

(i) Compute the optimal asymptotic constant

$$\alpha_k = \limsup_{g\to\infty} \frac{1}{g} \max\{r : \mathcal{J}_{g,1}(k) \text{ contains an undistorted } \mathbb{Z}^r\}.$$

(ii) Develop a nonabelian rate-distortion theory for residual fibres, using growth, hierarchical hyperbolicity, divergence, or asymptotic cones.

(iii) Compare information-theoretic lower bounds with explicit algorithms for choosing markings from period, nilpotent, or level data.

(iv) Extend the invisible-flat framework to character varieties, local systems, Higgs bundles, hyperkaehler moduli, and derived categories with large autoequivalence groups.

References

[BBF] M. Bestvina, K. Bromberg, and K. Fujiwara, *Constructing group actions on quasi-trees and applications to mapping class groups*, Publications mathematiques de l'IHES 122 (2015), 1–64.

[BM] J. Behrstock and Y. Minsky, *Dimension and rank for mapping class groups*, Annals of Mathematics 167 (2008), 1055–1077.

[BLM] J. S. Birman, A. Lubotzky, and J. McCarthy, *Abelian and solvable subgroups of the mapping class groups*, Duke Mathematical Journal 50 (1983), 1107–1120.

[CP] T. Church and A. Putman, *Generating the Johnson filtration*, Geometry & Topology 19 (2015), 2217–2255.

[FM] B. Farb and D. Margalit, *A Primer on Mapping Class Groups*, Princeton University Press, 2012.

[FLM] B. Farb, A. Lubotzky, and Y. Minsky, *Rank-one phenomena for mapping class groups*, Duke Mathematical Journal 106 (2001), 581–597.

[Joh] D. Johnson, *An abelian quotient of the mapping class group $\mathcal{I}_g$*, Mathematische Annalen 249 (1980), 225–242.

[MM] H. Masur and Y. Minsky, *Geometry of the complex of curves II: Hierarchical structure*, Geometric and Functional Analysis 10 (2000), 902–974.

[MP] J. Malestein and A. Putman, *Pseudo-Anosov dilatations and the Johnson filtration*, Groups, Geometry, and Dynamics 10 (2016), 1409–1433.

[Put] A. Putman, *The Johnson homomorphism and its kernel*, Journal fuer die reine und angewandte Mathematik 735 (2018), 109–141.

[Vau] W. R. Vautaw, *Abelian subgroups of the Torelli group*, Algebraic & Geometric Topology 2 (2002), 157–170.