A Toy Model of Feedback-Indexed Adaptation in Cross-System Qualitative Mapping

Higher-Order State Transitions over a Fixed Partial-Mapping Realization

Kostiantyn Osmolovskyi
Independent Researcher, Ukraine

Version 2.0, 2026

Abstract

This paper presents a provenance-bearing toy realization of feedback-indexed adaptation over a fixed partial cross-system qualitative-mapping substrate. The mapping substrate is inherited from a finite-coordinate realization of partial qualitative mapping [2]; the present model does not redefine its source projection, partial domain, target constructor, formal graph, quality profile, or scalarization. Instead, it adds a higher-order adaptation state with a mapping-active coordinate $\xi \in \mathbb{R}^3$ and a mapping-inactive latent coordinate, typed feedback records addressed to current state versions, a three-valued application assessment, positive adaptation transitions, and a fixed projection into the inherited mapping context. Feedback does not imply correction, and positive application does not imply state-value change, mapping-context change, established map change, or favorable fixed-profile quality change.

The simulation design specifies six deterministic feedback schedules, two 121×121 parameter surfaces, all 40,401 final records of each of four deterministic source families, and three separate Monte Carlo strata. The deterministic cases distinguish assessed-but-not-applied feedback, applied no-op, latent-only state change, mapping-active change, alternating mapping-active change, and positive adaptation on an inherited unmapped source. Because the inherited mapping domain is independent of ξ in this realization, domain membership remains invariant while higher-order states and local mapped assignments may change. The common S-family schedule is not an improvement operator: mean local loss rises from 0.0711 to 0.1050 in S1 and from approximately 0.046 to 0.234–0.245 in S2–S4. Monte Carlo mapped counts remain inherited from the fixed substrate, while stepwise local quality can move in either direction. These are outputs of the declared toy design and sampling laws, not empirical prevalence estimates.

Keywords: feedback; adaptation; cross-system mapping; higher-order state; partial mapping; provenance; definedness; mapping quality; toy model; cognitive architecture.

1 Analytic Target and Architectural Boundary

The present model asks a narrow question: what additional structure is needed to represent feedback-indexed adaptation without collapsing feedback, application, higher-order state change, mapping-context change, map change, domain change, and local quality change into one scalar trajectory? The answer is implemented as a higher-order layer over a finite-coordinate partial-mapping realization [2, 3, 4, 5].

The constitutive chain is

$$
\begin{aligned}
&\text{fixed source/receiver mapping anchor} \\
&\quad + \ A^{\mathrm{ad}}_{AB} + E^{\mathrm{fb}}_{AB} + C^{\mathrm{ad}}_{AB} \\
&\quad \longrightarrow \text{AdApply} \in \{\text{undefined}, 0, 1\} \\
&\quad \longrightarrow T^{\mathrm{ad}}_{AB} \text{ only when } \ \text{AdApply} = 1 \\
&\quad \longrightarrow \text{next higher-order adaptation state} \\
&\quad \longrightarrow \pi^{\mathrm{map}}_{AB} \\
&\quad \longrightarrow \text{inherited } \chi_{AB} \\
&\quad \longrightarrow \text{inherited partial mapping and domain} \\
&\quad \longrightarrow \text{fixed-profile quality readout.}
\end{aligned}
$$

The model therefore does not construct a second mapping rule. It does not add a separate mapping-grounding layer, reverse mapping, multi-system aggregation, or general trajectory-dependent regulation. It also does not use a scalar readability/misreadability complement.

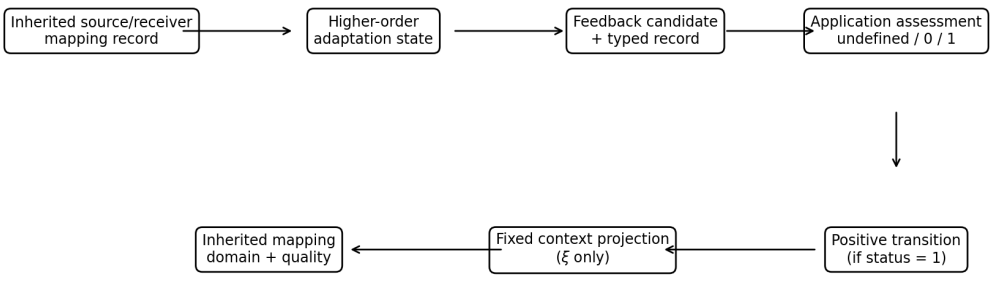


Figure 1: Feedback-indexed adaptation architecture over an inherited partial-mapping realization. The latent state coordinate is deliberately excluded from the fixed projection into the mapping context.

2 Inherited Mapping Substrate

Every process is anchored to one source record and one receiver-context record from the finite-coordinate mapping realization [2]. Those records determine the inherited source-side qualitative application, reduced source coordinate, receiver context, partial mapping family, domain-support rule, target construction, formal graph application, and fixed quality profile. The adaptation layer treats those mapping definitions as fixed inputs.

The inherited mapping context contains the receiver-side coordinates and a three-dimensional perturbation ξ. Crucially, the inherited domain-support rule does not depend on ξ. Hence, for a fixed source and fixed receiver context,

$$\mathrm{Dom}(\phi^{\chi_k}_{AB}) = \mathrm{Dom}(\phi^{\chi_0}_{AB})$$

for every adaptation state projected by the rule below. This is a realization-specific consequence of the selected mapping substrate. It is not a general theorem that adaptation cannot change mapping-readability domains [4].

The inherited quality scalar $\ell_{AB}^{\mathrm{coord}}$ is evaluated only under the fixed profile and scalarization of the finite-coordinate realization [2]. A lower value is locally more favorable under that frame. It does not constitute mapping readability, understanding, truth, semantic correction, or empirical success.

3 Higher-Order Adaptation State

Each adaptation process binds one source anchor and one receiver-context anchor. Its adaptation state is

$$A_{AB,k}^{\mathrm{ad}} = (\mathrm{process}, \mathrm{stateVersion}_k, \xi_k, \lambda_k),$$

where $\xi_k \in \mathbb{R}^3$ is mapping-active and $\lambda_k \in \mathbb{R}$ is a mapping-inactive latent coordinate. This higher-order state follows the structural separation formalized for feedback-indexed adaptation [5]. Attempt index and state version are distinct. Attempt index advances on every feedback attempt; state version advances only after a positive application.

The fixed projection is

$$\pi_{AB}^{\mathrm{map}}(A_{AB,k}^{\mathrm{ad}}) = \chi_{AB}(\text{exact receiver context}, \xi_k).$$

The latent coordinate λ_k is excluded. Consequently, a higher-order state may change while the projected mapping context remains byte-for-byte identical. Conversely, a change in ξ_k changes the context record but does not by itself prove extensional map change.

4 Feedback, Application, and Transition

A typed feedback record contains process identity, channel-event identity, attempt index, addressed state identity/version, an admission signal $a \in [0, 1]$, a proposed $\Delta\xi \in \mathbb{R}^3$, and a proposed $\Delta\lambda \in \mathbb{R}$. These are neutral synthetic control payloads. The record does not self-classify as corrective, misleading, truthful, successful, or effective.

The adaptation context uses

$$\tau_{\mathrm{app}} = 0.40, \qquad g_\xi = g_\lambda = 0.50.$$

Application is undefined when the typed process/state/feedback triple cannot be fully assessed. When assessment is defined,

$$\mathrm{AdApply} = \begin{cases} 0, & a < 0.40, \\ 1, & a \geq 0.40. \end{cases}$$

Only $\mathrm{AdApply} = 1$ licenses a transition:

$$\xi_{k+1} = \xi_k + 0.50\,\Delta\xi_k, \qquad \lambda_{k+1} = \lambda_k + 0.50\,\Delta\lambda_k.$$

No hidden clipping is applied at the higher-order adaptation-state layer. If $\mathrm{AdApply} = 0$, no transition record is created and the current state is carried forward. An applied no-op is different: $\mathrm{AdApply} = 1$ with zero proposals creates a successor state version whose numerical state values are unchanged. Thus $\mathrm{AdApply} = 0$ is not an applied no-op, and undefined application is not application status zero.

5 Map, Domain, Image, and Quality Readouts

For each positive transition, the pre/post higher-order states are projected into the inherited mapping contexts and evaluated on the same exact source point. Four readouts are kept separate.

First, state-value change compares (ξ, λ) exactly. Second, context change compares the inherited χ records exactly. Third, map relation is typed as `SAME`, `CHANGED`, or `UNRESOLVED`. Exact identical contexts license `SAME`. A changed context plus a changed inherited target assignment for the common mapped source provides a witness for `CHANGED`. A changed context without such a witness is `UNRESOLVED`; absence of a witness does not prove map equality.

Fourth, full image status is `NOT_EVALUATED`. One selected source point cannot establish change of the full map image. The local quality relation is evaluated only when both pre/post inherited quality records exist under the exact same profile and scalarization. It is then `MORE_FAVORABLE`, `LESS_FAVORABLE`, or `EQUIVALENT`; otherwise it is `INCOMPARABLE_OR_UNDEFINED`. Loss reduction does not imply mapping-domain expansion or a general improvement in readability [4].

6 Deterministic Cases

The deterministic design uses six 20-attempt schedules. D0 uses a subthreshold admission signal; D1 applies a zero proposal; D2 changes only the latent coordinate; D3 repeatedly changes the mapping-active coordinate; D4 alternates opposite mapping-active proposals; and D5 applies mapping-active feedback to a source that is initially outside the inherited map domain.

Table 1: Deterministic feedback-adaptation schedules and realized status counts.

Scenario	base	app	no-app	no-op	state Δ	context Δ	map wit.	map unres.	final map	final ℓ
D0_assessed_not_applied	supportive_context	0	20	0	0	0	0	0	1	0.0176
D1_applied_noop	supportive_context	20	0	20	0	0	0	0	1	0.0176
D2_latent_only	supportive_context	20	0	0	20	0	0	0	1	0.0176
D3_mapping_shift	supportive_context	20	0	0	20	20	20	0	1	0.3176
D4_alternating_mapping	supportive_context	20	0	0	20	20	20	0	1	0.0176
D5_unmapped_active	receiver_pressure_context	20	0	0	20	20	0	20	0	–

The cases separate relations that a scalar feedback model would tend to collapse. D0 contains 20 defined assessments but no positive applications and remains exactly constant. D1 contains 20 positive applications and 20 successor state versions while producing no numerical state-value change. D2 changes the higher-order state on every application while leaving χ and the selected-source map assignment unchanged. D3 produces 20 mapping-context changes and 20 assignment-change witnesses; its local loss changes from 0.0176 to 0.3176. D4 also produces a map-change witness at every applied step, but the alternating proposals return the final ξ and final local loss to their initial values. Endpoint equality therefore does not erase intervening map changes. D5 remains outside the inherited domain on every step: its context changes, but there is no mapped selected-source assignment from which a map-change witness could be established, so those relations remain unresolved and quality remains undefined.

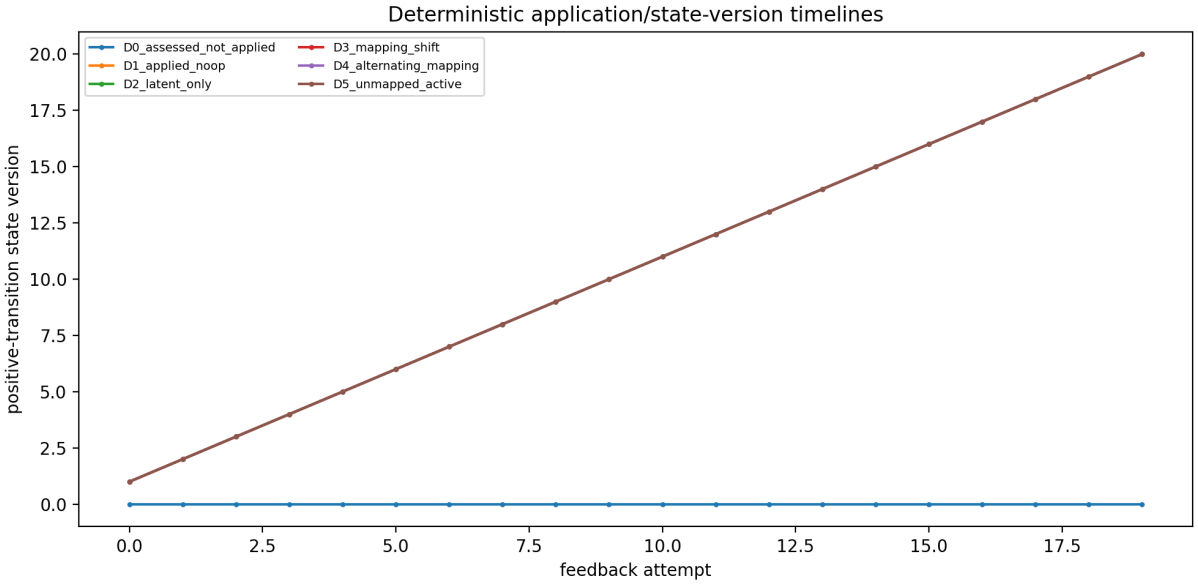


Figure 2: D0–D5 attempt/application/state/context/map-witness timelines. The panels expose assessed-but-not-applied attempts, applied no-op, latent-only state change, mapping-active change, alternating change, and active adaptation on an inherited unmapped source.

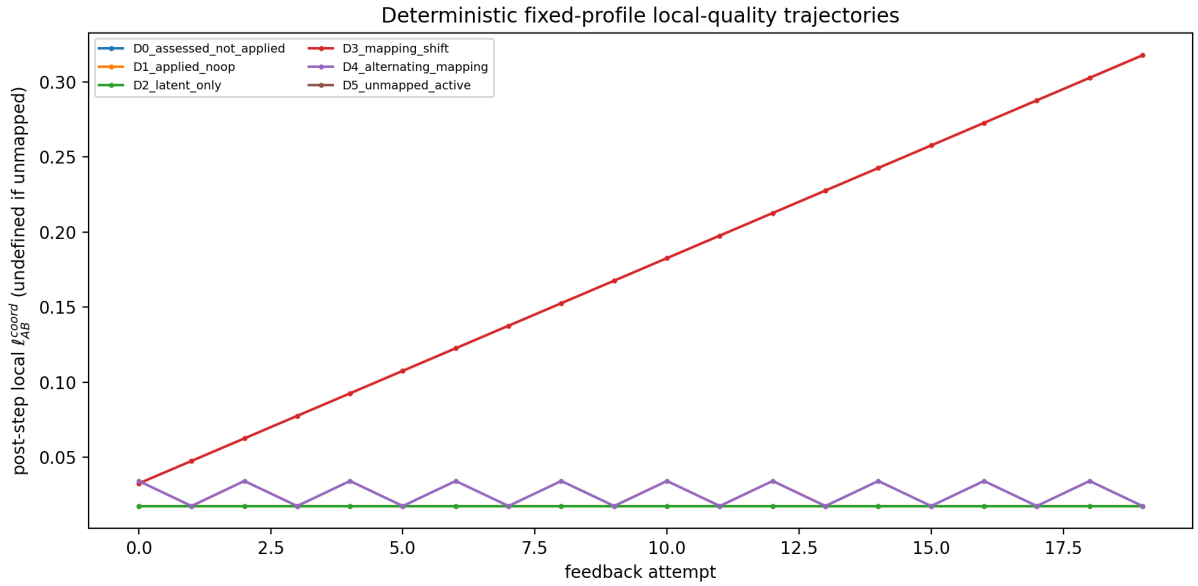


Figure 3: D0–D5 selected-source local $\ell_{AB}^{\mathrm{coord}}$ trajectories where fixed-profile quality is defined. D5 remains unmapped, so no numerical loss is fabricated. A lower value is only locally more favorable under the fixed frame.

7 Feedback Surfaces

The first parameter surface varies admission signal and the first mapping-active proposal coordinate over a 121×121 grid. The second surface varies the first mapping-active proposal and the mapping-inactive latent proposal, again over 121×121 cells. Both use the same supportive mapping anchor and 20 attempts per cell.

For G1, all 14,641 final cells remain mapped because domain membership is independent of

ξ. Positive numerical state/context change occurs in 8,760 cells. Every such cell has at least one selected-source assignment-change witness, while 4,307 cells also contain at least one later changed-context step for which no new assignment-change witness is available, primarily after coordinate clipping makes the selected target assignment stationary. Final local loss ranges from 0.0109 to 0.2842. Final-two state/context constancy occurs in 5,881 cells.

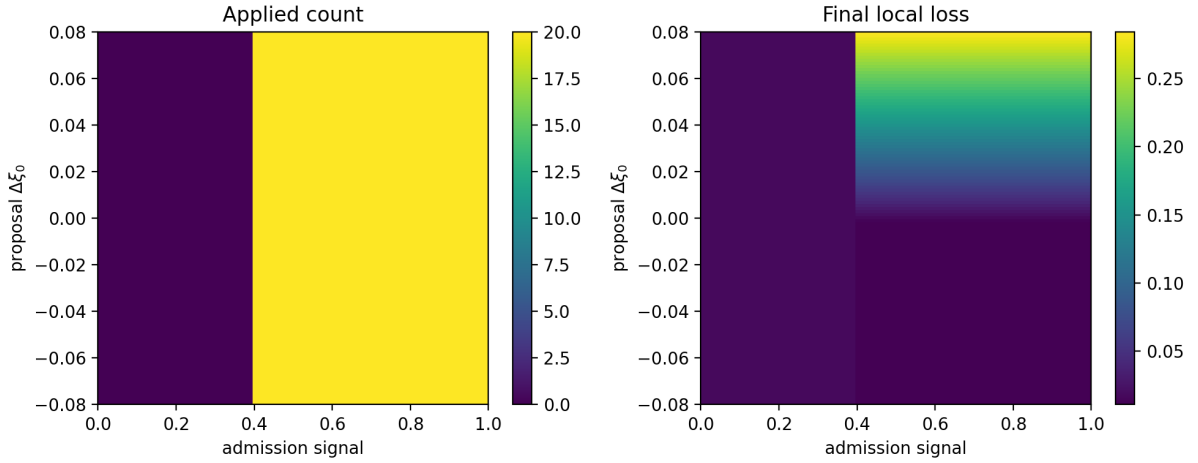


Figure 4: G1: admission signal by mapping-active proposal $\Delta\xi_0$. The application threshold controls whether transitions occur, while repeated ξ updates can eventually stop changing the selected target because of inherited target-coordinate clipping.

For G2, all 14,641 cells are positively applied on every attempt. State values change in 14,640 cells, whereas the projected mapping context changes in 14,520 cells; the 120 cells with $\Delta\xi_0 = 0$ but nonzero latent proposal therefore provide a surface-level state-change-without-context-change region. Only the single zero/zero proposal cell is final-two state-constant, while all 121 cells with $\Delta\xi_0 = 0$ are context-constant. Final local loss again ranges from 0.0109 to 0.2842.

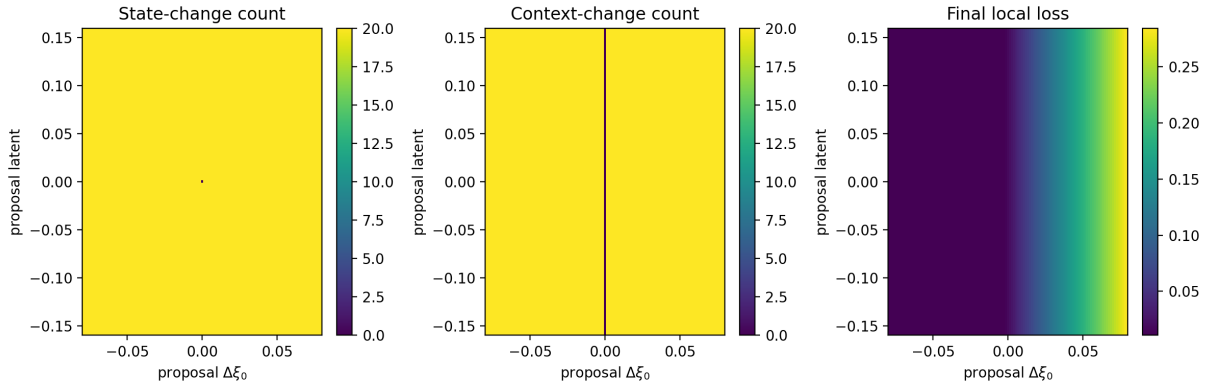


Figure 5: G2: mapping-active proposal $\Delta\xi_0$ by mapping-inactive latent proposal $\Delta\lambda$. The latent axis changes higher-order state without changing the projected mapping context.

8 Complete S1–S4 Families

All 40,401 final records of each S1–S4 source family are used once in their declared order. No new source sampling occurs. Each record receives the same 20-attempt schedule with admission 0.80, $\Delta\xi = (0.03, -0.02, 0.01)$, and $\Delta\lambda = 0.02$.

Table 2: Complete S1–S4 source families under the common feedback schedule.

Fam.	N	applied	map wit.	map unres.	final map	mean init. ℓ	mean final ℓ	state const.
S1	40401	808020	808020	0	40401	0.0711	0.1050	0
S2	40401	808020	808020	0	40401	0.0467	0.2378	0
S3	40401	808020	808020	0	40401	0.0457	0.2446	0
S4	40401	808020	808020	0	40401	0.0468	0.2340	0

Every S-family source is mapped both before and after the feedback schedule, and domain membership remains invariant. Every attempt is positively applied and changes the mapping-active state. The selected-source assignment changes on every step for these ranges. Exact final-two state and context constancy therefore occur in zero rows, while domain constancy occurs in all 40,401 rows of every family.

The local scalar results are deliberately not tuned toward improvement. Mean $\ell_{AB}^{\mathrm{coord}}$ changes from 0.0711 to 0.1050 in S1, from 0.0467 to 0.2378 in S2, from 0.0457 to 0.2446 in S3, and from 0.0468 to 0.2340 in S4. These shifts under the selected fixed profile are outputs of the declared realization. They demonstrate that feedback-indexed application is not, by definition, a quality-improvement operator.

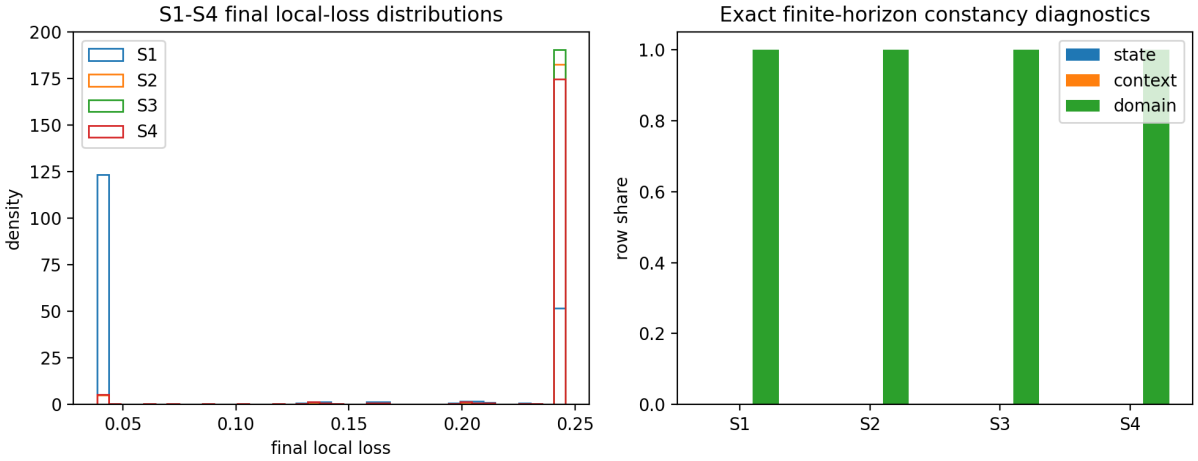


Figure 6: S1–S4 final local-loss distributions and final-two constancy diagnostics. Domain constancy is complete by construction of this realization; state/context constancy is absent under the continuously active fixed schedule.

9 Monte Carlo Strata

Three Monte Carlo final-record strata are retained separately: MC-M1 ($N = 9,584$), MC-M2 ($N = 14,416$), and MC-M3 ($N = 16,948$). The adaptation layer adds only feedback draws over 20 attempts. The seeds are 2051, 2052, and 2053. For each stratum the draw order is fixed as the full admission matrix, the full three-coordinate proposal tensor, and the full latent-proposal matrix.

Table 3: Monte Carlo strata. Strata are reported separately; counts characterize the declared sampling design only.

Fam.	N	applied	no-app	map wit.	map unres.	final map	mean final ℓ	1/map	1/unmap
MC-M1	9584	115036	76644	93480	21556	7782	0.0711	7782	1533
MC-M2	14416	172428	115892	140606	31822	11752	0.0706	11752	2260
MC-M3	16948	203364	135596	167775	35589	13972	0.0711	13972	2528

Mapped counts remain exactly 7,782, 11,752, and 13,972, respectively, because the inherited domain is invariant to the mapping-active coordinate in this realization. The corresponding unmapped counts are 1,802, 2,664, and 2,976. Mean mapped local loss changes from approximately 0.056–0.057 initially to approximately 0.071 finally in all three strata. Within mapped rows, both more-favorable and less-favorable individual steps occur. Final-two state/context constancy occurs in 3,835, 5,901, and 6,810 rows, mainly where the last attempt is assessed but not applied. These are finite-horizon constancy diagnostics, not claims of metric convergence.

The inherited pre-readout status remains separate from mapping-domain membership [1]. For MC-M1, for example, 7,782 rows have positive upstream pre-readout and are mapped, while 1,533 have positive upstream pre-readout but remain unmapped; 269 have zero pre-readout and remain unmapped. No feedback step changes those domain memberships in this realization.

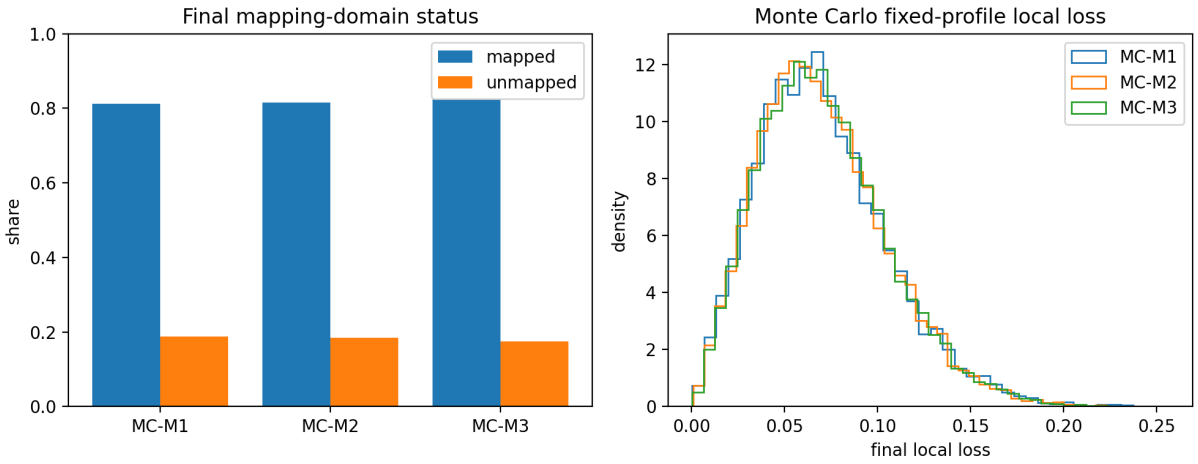


Figure 7: Monte Carlo final mapping-domain status and fixed-profile local-loss distributions. Strata are kept separate. Frequencies are design/sampling-law outputs and are not empirical prevalence estimates.

10 Definedness and Stabilization Boundaries

The implementation keeps distinct denominators for candidate availability, application definedness, assessed-not-applied attempts, positive applications, applied no-op, state change, context change, map-change witness, unresolved map relation, domain equality, image non-evaluation, pre/post quality evaluability, and local quality direction. In the standard runs all feedback records are synthetically complete, so application is always defined; the undefined route nevertheless remains executable and is tested separately.

Several non-equivalences are therefore mandatory:

- feedback availability $\neq$ positive application;
- positive application $\neq$ numerical state-value change;
- state change $\neq$ mapping-context change;
- mapping-context change $\neq$ proven extensional map change;
- absence of a selected-source witness under changed context $\neq$ proof that the maps are equal;
- one selected-source assignment $\neq$ the full map image;
- unmapped $\neq$ distorted or maximal loss;

- local loss direction $\neq$ mapping-domain or general readability change;
- feedback-indexed adaptation $\neq$ correction, understanding, or truth convergence.

The stabilization diagnostics are finite-horizon eventual-constancy indicators operationalized by final-two equality of state values, projected contexts, established map relation where available, and domain membership. They are not generic metric-convergence claims. Because the domain is invariant under the selected projection, domain constancy is structurally expected here and must not be generalized beyond this realization.

11 Computational Reproducibility

The public computational supplement contains the adaptation-layer equations, compact deterministic and Monte Carlo result records, the two parameter-surface arrays, and a reproduction script. The script regenerates all figures and tables and checks the declared application threshold, transition gains, no-op and latent-only cases, deterministic schedules, surface counts, source-family summaries, Monte Carlo sample sizes and seeds, mapping-domain constancy, undefined-quality handling, and denominator partitions. These checks concern the scientific realization and reported outputs; they do not encode document-version history or package-integrity bookkeeping.

12 Limitations

This is a deliberately narrow toy realization. Feedback payloads are exogenous synthetic records rather than endogenous observations generated by a receiving system. The latent coordinate has no direct mapping effect except through higher-order state bookkeeping. The inherited domain cannot change because the selected mapping-domain rule does not depend on the mapping-active coordinate. Map change is established only through an assignment-change witness at the selected source; full extensional map equality or image change is not enumerated. The local $\ell_{AB}^{\mathrm{coord}}$ comparison uses one fixed profile and scalarization. No empirical interaction, reverse-direction mapping, multi-system aggregation, or general trajectory-dependent regulation is constructed.

The Monte Carlo frequencies are not empirical prevalence. They describe only the declared synthetic sampling laws and inherited record strata. Likewise, neither lower local loss nor eventual constancy is interpreted as successful communication, semantic correction, truth, or convergence of experience.

13 Conclusion

The model implements a typed higher-order adaptation layer over a fixed partial-mapping substrate. It makes four distinctions executable: feedback assessment versus application, application versus state-value change, state change versus projected mapping-context change, and context change versus established map change. Under the declared projection, adaptation can alter higher-order state and local mapped assignments while leaving mapping-domain membership invariant. The common S-family schedule also frequently makes the selected fixed-profile loss less favorable. Feedback is therefore not correction by definition, and application is not improvement by definition. The purpose of the realization is to preserve these architectural separations while providing a finite computational surface on which they can be inspected.

References

[1] Osmolovskyi, K. (2026). *A Toy Model of Cross-System Pre-Readout Constraint under Identity-Continuation Variation.* Zenodo. https://doi.org/10.5281/zenodo.20189695

[2] Osmolovskyi, K. (2026). *A Finite-Coordinate Toy Realization of Partial Cross-System Qualitative Mapping.* Zenodo. https://doi.org/10.5281/zenodo.20189802

[3] Osmolovskyi, K. (2026). *Partial Cross-System Qualitative Mapping in Cognitive Systems.* Zenodo. https://doi.org/10.5281/zenodo.20190149

[4] Osmolovskyi, K. (2026). *Readability Domains in Cross-System Qualitative Mapping.* Zenodo. https://doi.org/10.5281/zenodo.20190329

[5] Osmolovskyi, K. (2026). *Adaptive Cross-System Qualitative Mapping under Feedback.* Zenodo. https://doi.org/10.5281/zenodo.20190862