

Breaking the Loop

Iteration Drift in Probabilistic Systems

(v1.0)

Large language models generate coherent text with high fluency. They are less reliable at preserving structural integrity across long iterative edits.

In extended editing sessions, a recurring pattern appears:

A revision improves a paragraph.

A follow-up correction adjusts wording.

After several rounds, a section is missing, an earlier clarification disappears, or older phrasing reappears.

A restart is requested.

The system produces a new version.

That new version is anchored to the current visible state, not the original structure.

This behavior is a property of iterative conditional generation.

Convergence in Rule-Driven Tools

In rule-driven tools such as compilers or code formatters, corrections converge.

Run a formatter once — indentation is adjusted.

Run it again — no further changes occur.

Once the artifact satisfies the rules, additional runs produce no structural modification.

These tools are often idempotent: applying the transformation repeatedly yields the same result after the first correction.

Generative systems do not inherently behave this way.

When asked to “improve this section,” the model does not apply a fixed rule. It generates a new version conditioned on the visible context. If the context has already drifted, the output reflects that drift.

Repeated correction does not guarantee convergence toward a stable structure.

Mutation vs Reconstruction

Two interaction patterns can be distinguished.

Mutation is incremental editing within the current visible state. The system modifies existing text while conditioning on whatever remains in context.

Reconstruction re-anchors generation to an explicit structure — by restating the outline, reinjecting a canonical draft, or regenerating from a defined reference.

Mutation modifies the present state.

Reconstruction regenerates against a stable reference.

If the reference is no longer visible, mutation compounds structural drift.

These are interaction-level distinctions, not internal architectural modes.

Analog Drift

In analog media duplication, each copy slightly degraded fidelity. The first copy remained close to the source. Subsequent generations accumulated noise.

Each copy was usable. None restored the original.

Iterative generative edits behave similarly when each pass conditions on the most recent state rather than a stable reference.

Mechanism

Language models generate text by predicting continuations based on visible context. Semantic relationships and positional structure influence output.

Text that remains in context shapes subsequent generation. Text that is absent does not.

If a canonical structure falls out of view, subsequent generations reconstruct from what remains.

Practical Implication

Treating every issue as a local patch maintains mutation mode.

Reintroducing explicit structure — restating the outline, reinjecting the full draft, or regenerating from a defined reference — restores a stable anchor for generation.

This compensates for non-idempotent behavior in iterative workflows.