# Textual Equilibrium Propagation (TEP)

## Local-first optimization for deep compound AI systems.

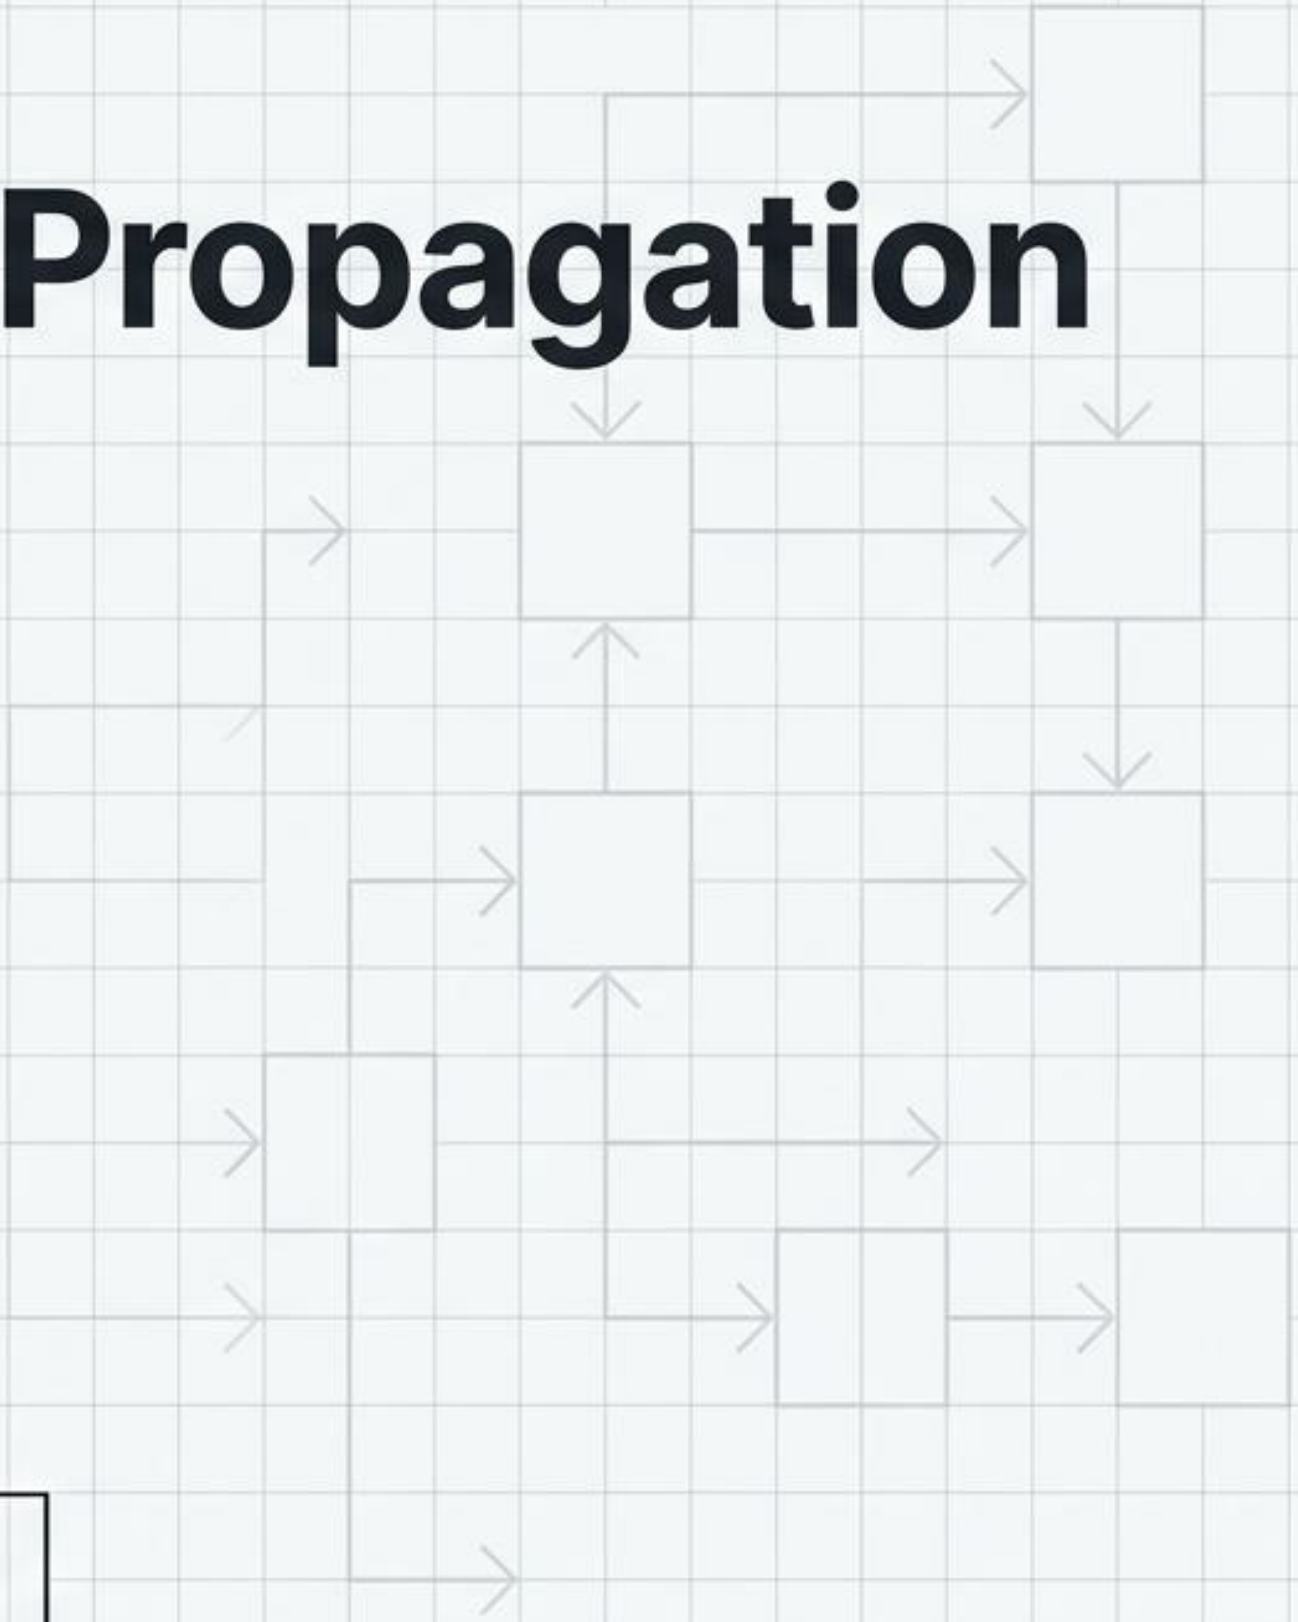

Minghui Chen, Wenlong Deng, James Zou, Han Yu, Xiaoxiao Li.

# Multi-step agentic workflows require long-horizon coordination.

## The Past: Single-Agent Architecture

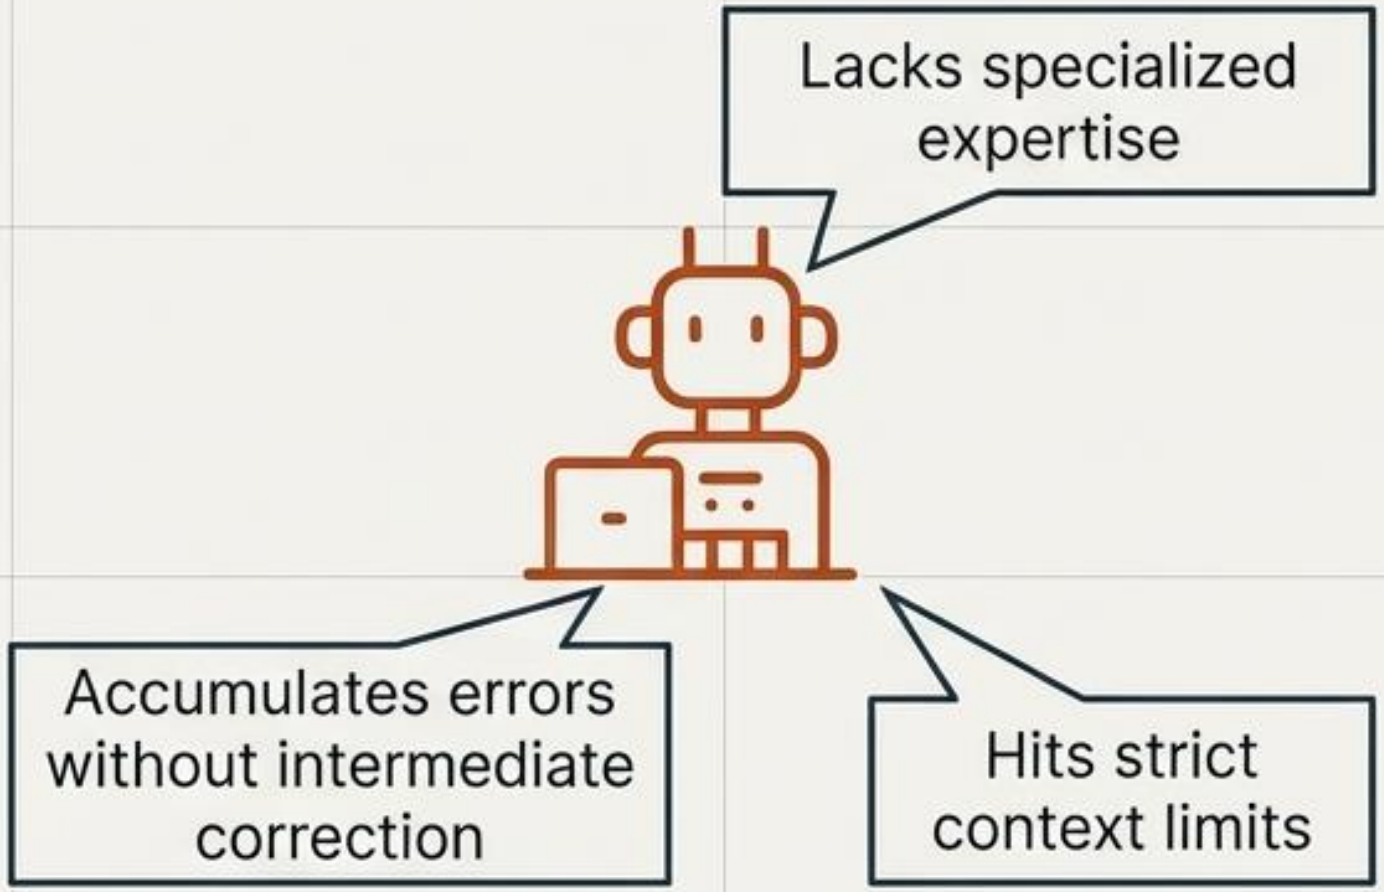


## The Present: Deep Compound AI

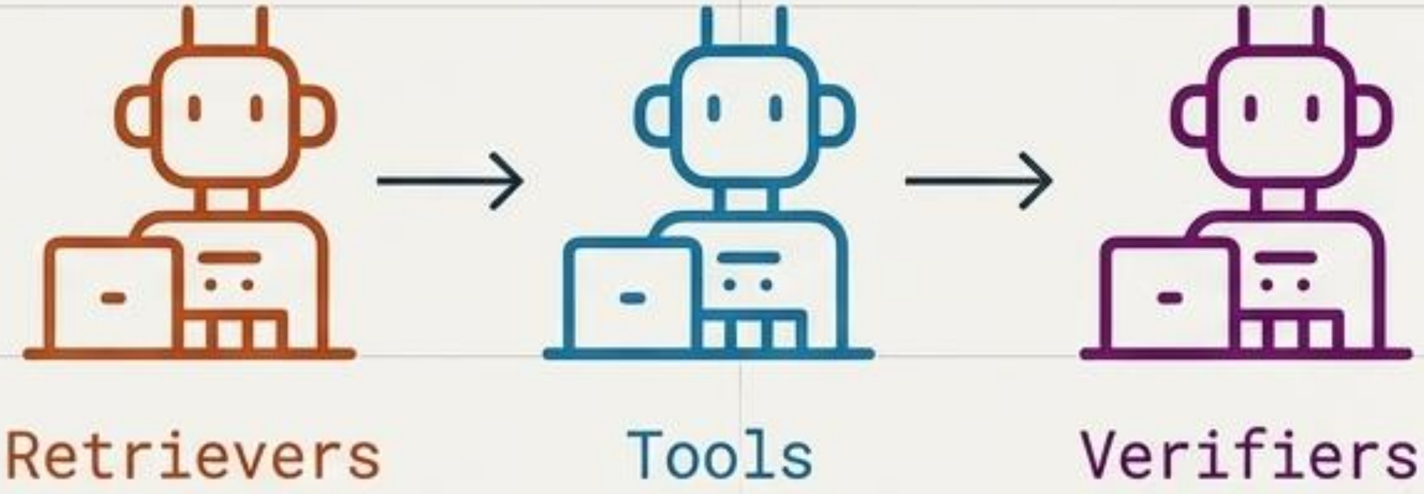


**Insight:** As task complexity grows, we no longer optimize **single models** — we **must optimize entire execution graphs.**

# 'Automatic differentiation via text' enables pipeline optimization.

Baseline Mechanism: TextGrad (Stochastic Computation Graphs)

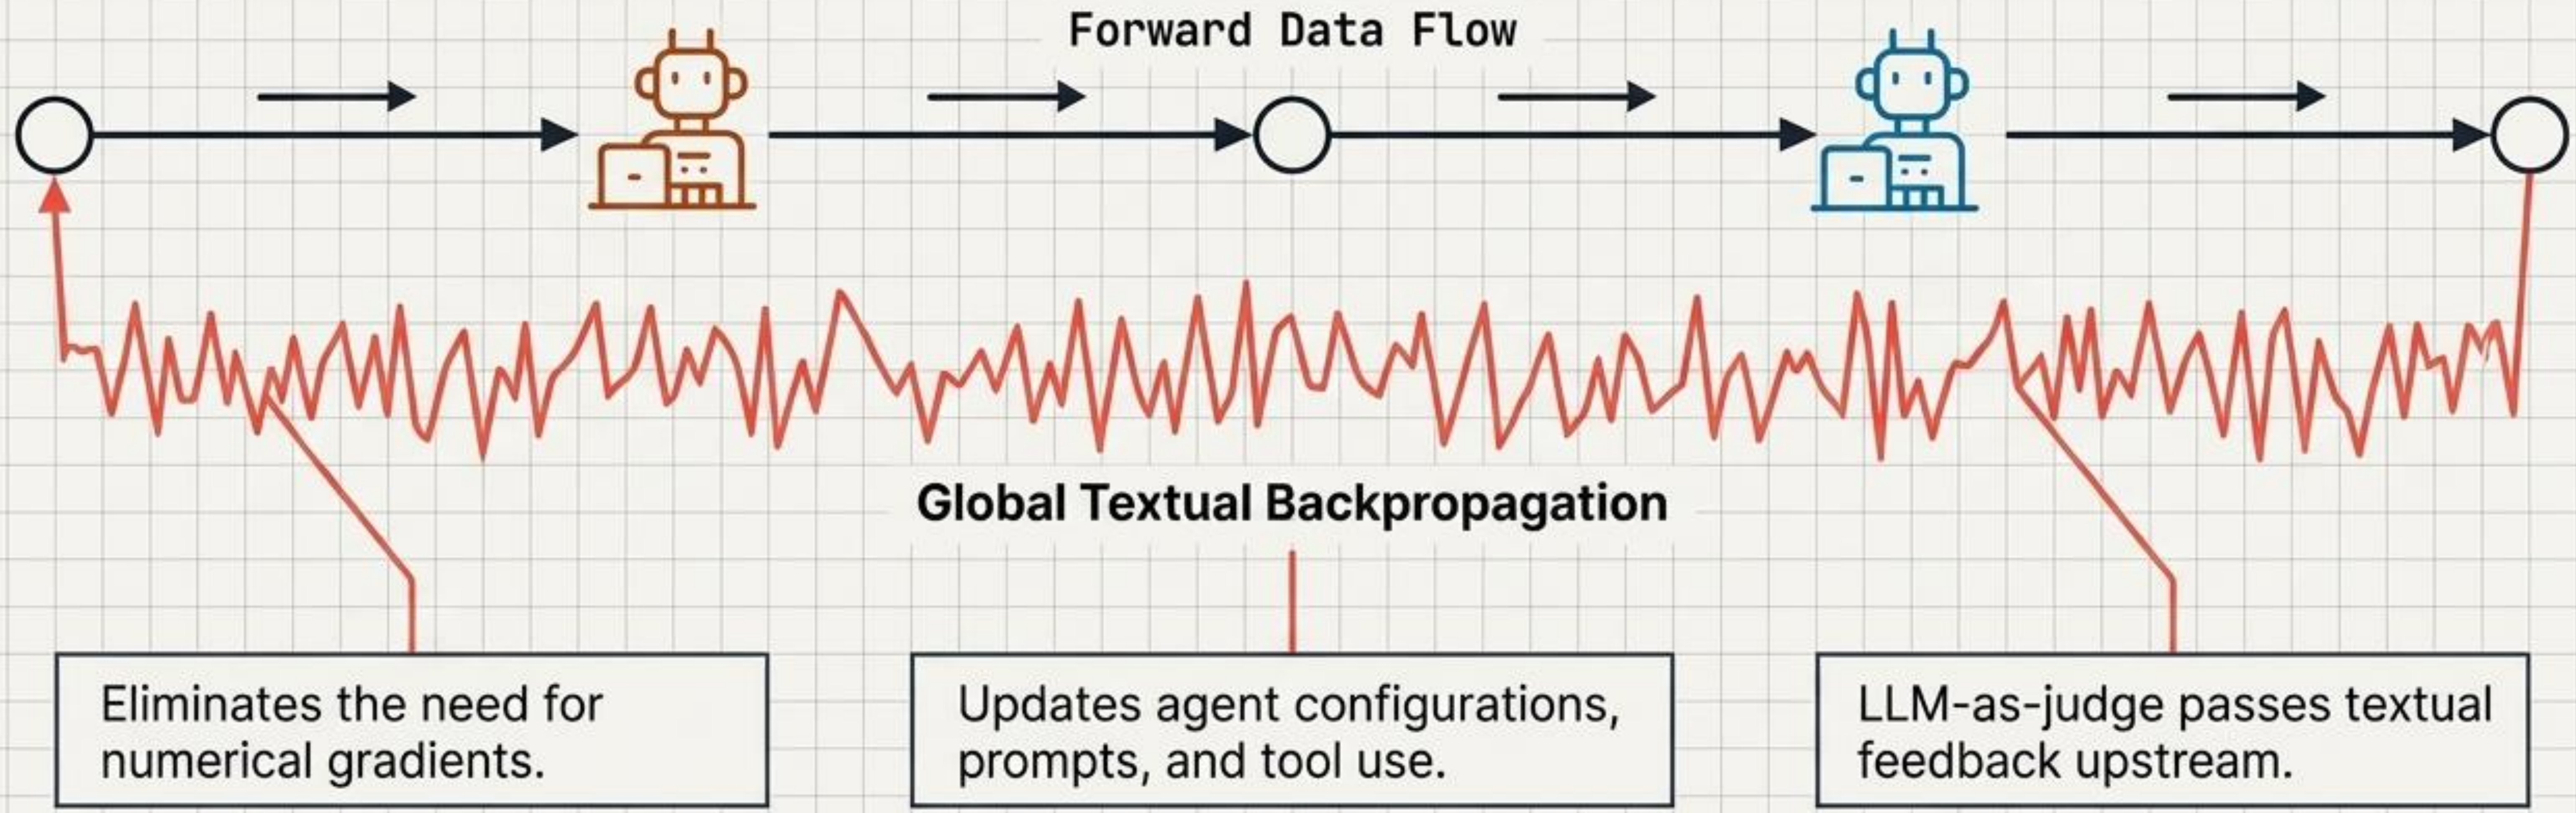


Takeaway: Highly effective for shallow pipelines, but fundamentally unstable across long execution distances.

# Compound AI systems have hit the exact same depth-scaling wall.

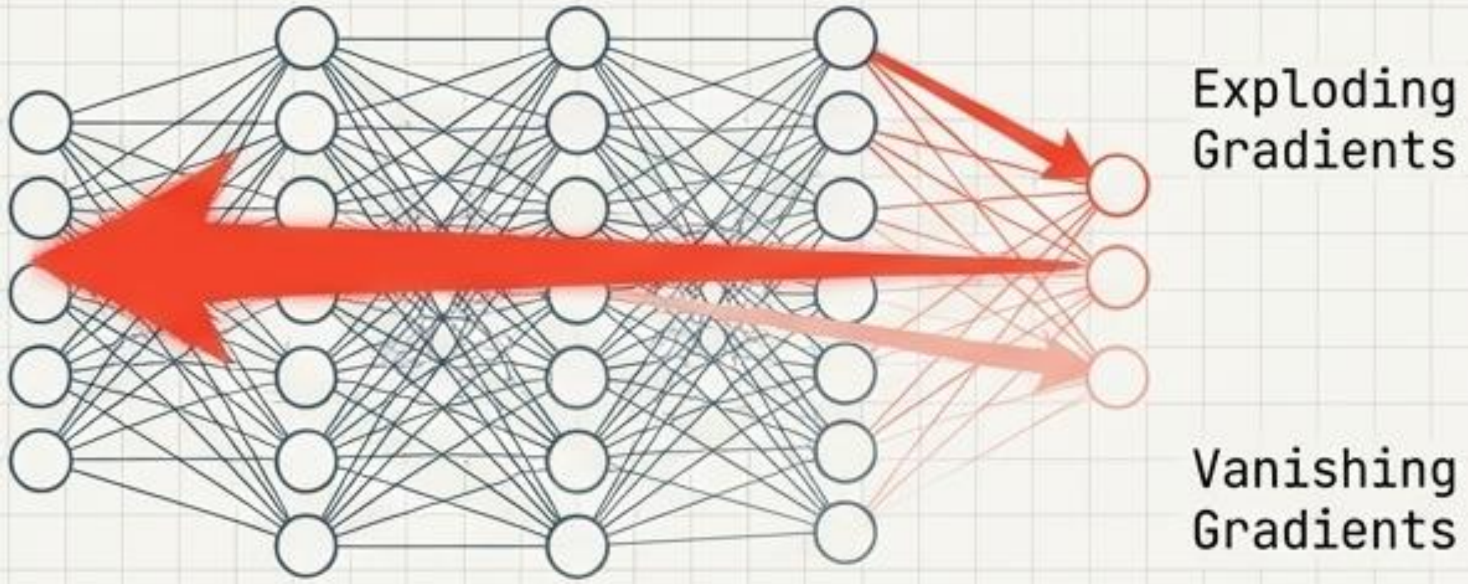


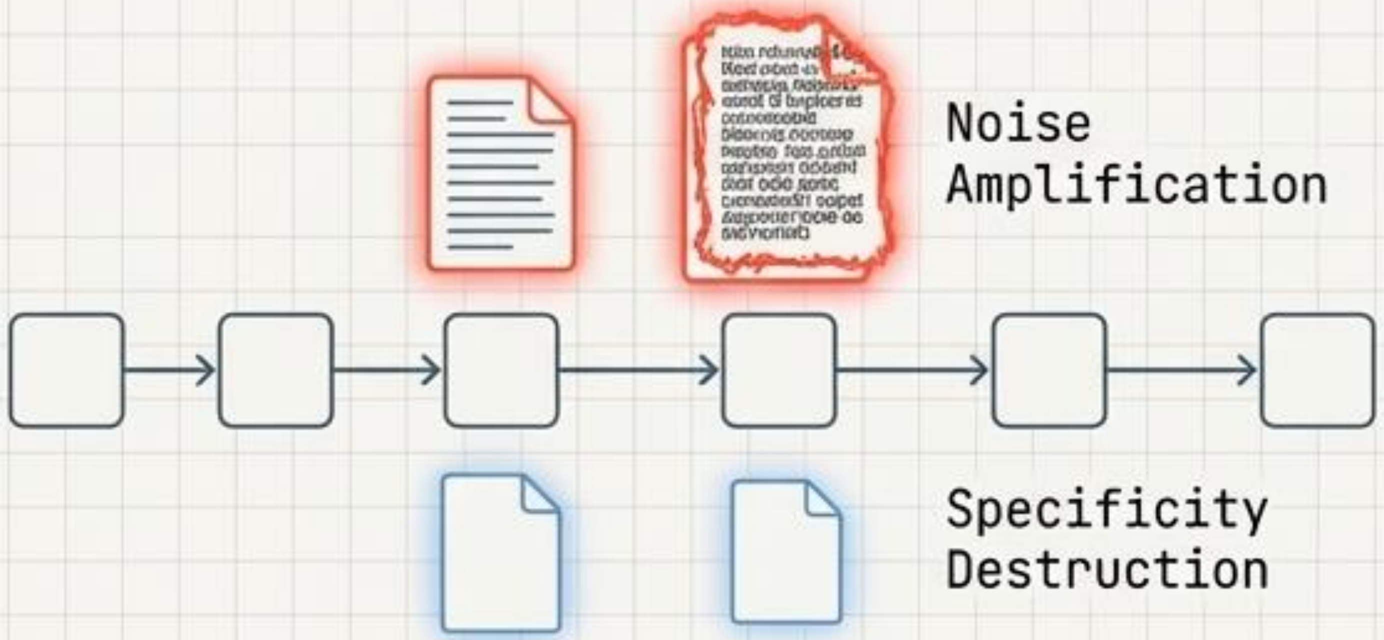


1990s: Biological implausibility of numerical backpropagation (Bengio et al., 1994). **Gradients vanish or explode at depth.**

Gradients either amplify exponentially or decay to zero during backpropagation, making deep network training unstable.

**Analogous Signal Degradation:**

Each agent layer distorts the optimization signal, rendering global backwards chains ineffective for deep workflows.

Today: **Contextual impossibility of global textual backpropagation.** Textual feedback exponentially amplifies noise or destroys specificity.

Textual feedback passed upstream either introduces compounding errors or becomes overly generic, failing to provide precise optimization signals.

# Tracking a single optimization signal across four hops.

We trace how global backpropagation distorts a precise fix moving upstream.

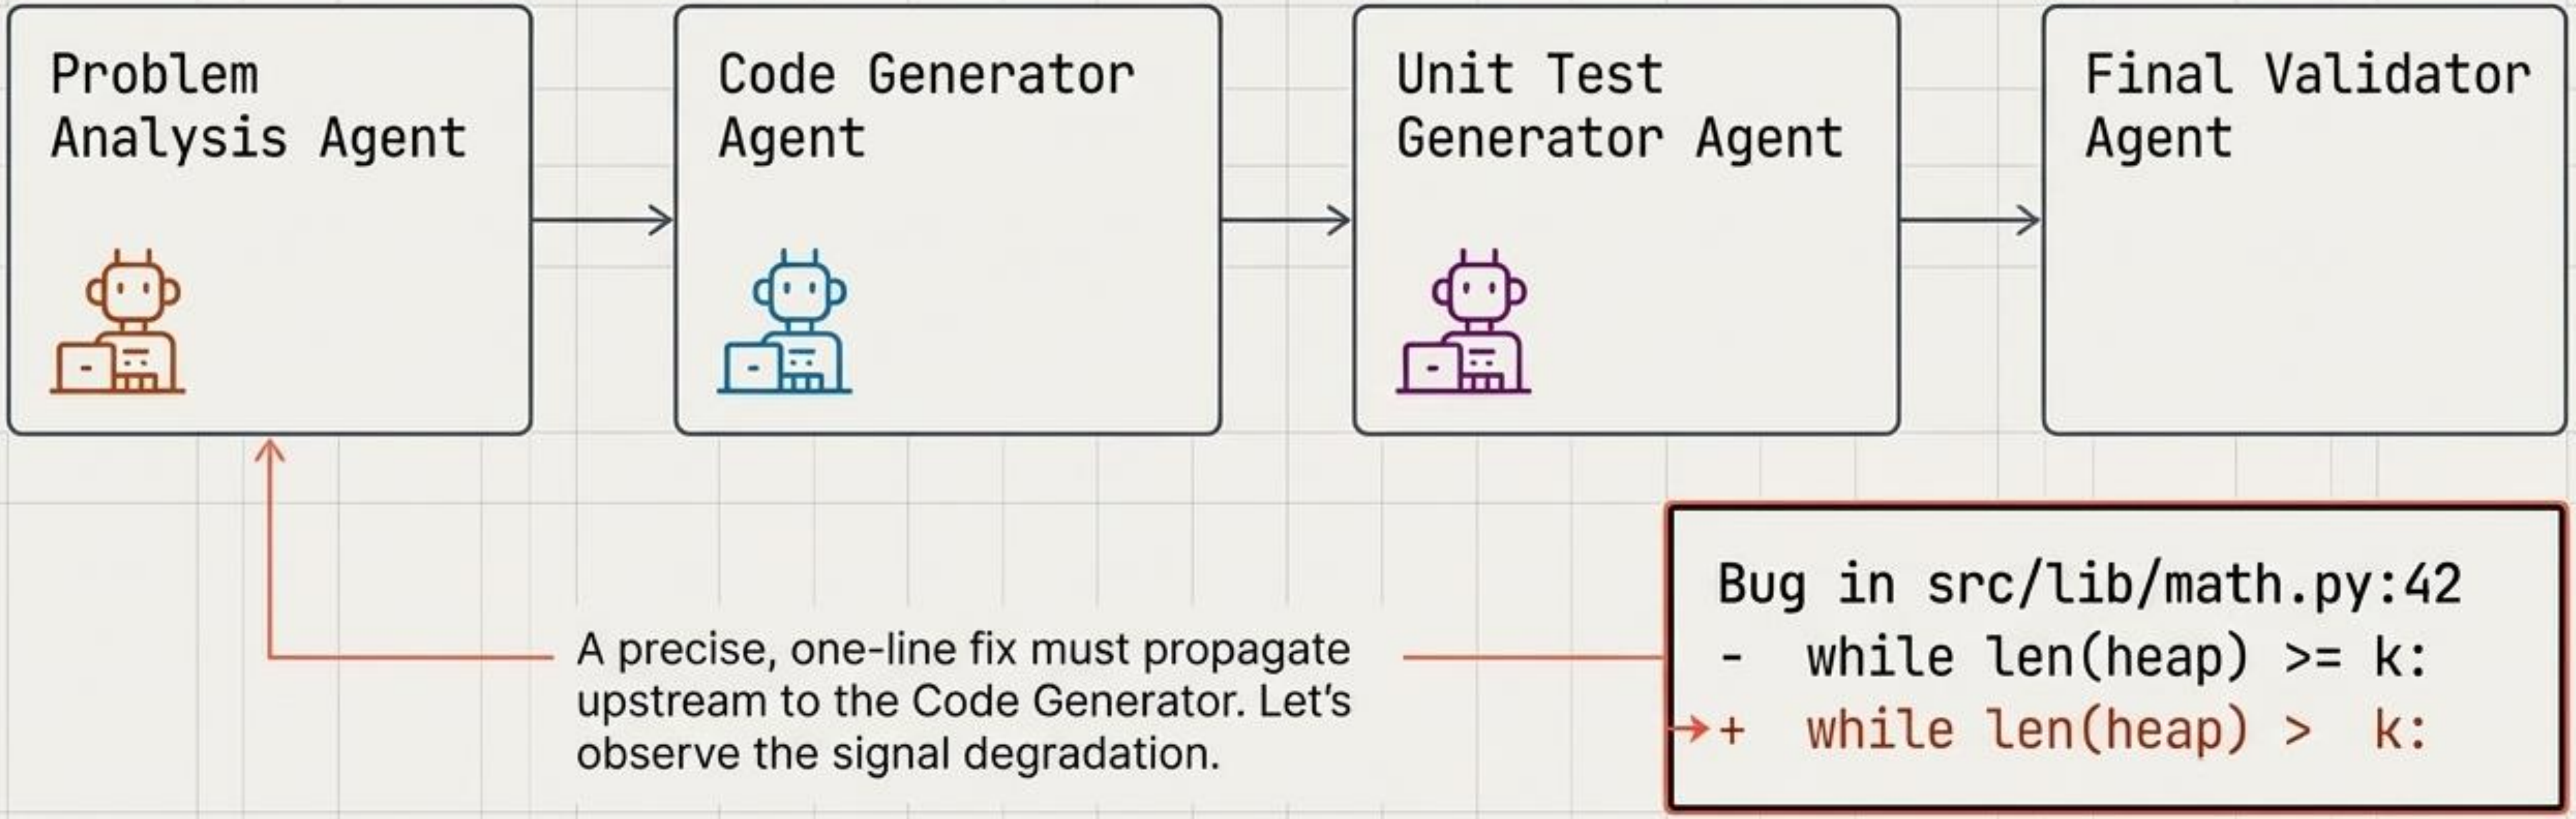

# Pathology 1: The Exploding Textual Gradient

**Symptom:** Feedback accumulates across layers.

**Mechanism:** Each agent adds context to preserve information, causing exponential token growth.

**Result:** The actionable fix (> vs >=) is buried beneath accumulated text ("lost in the middle").

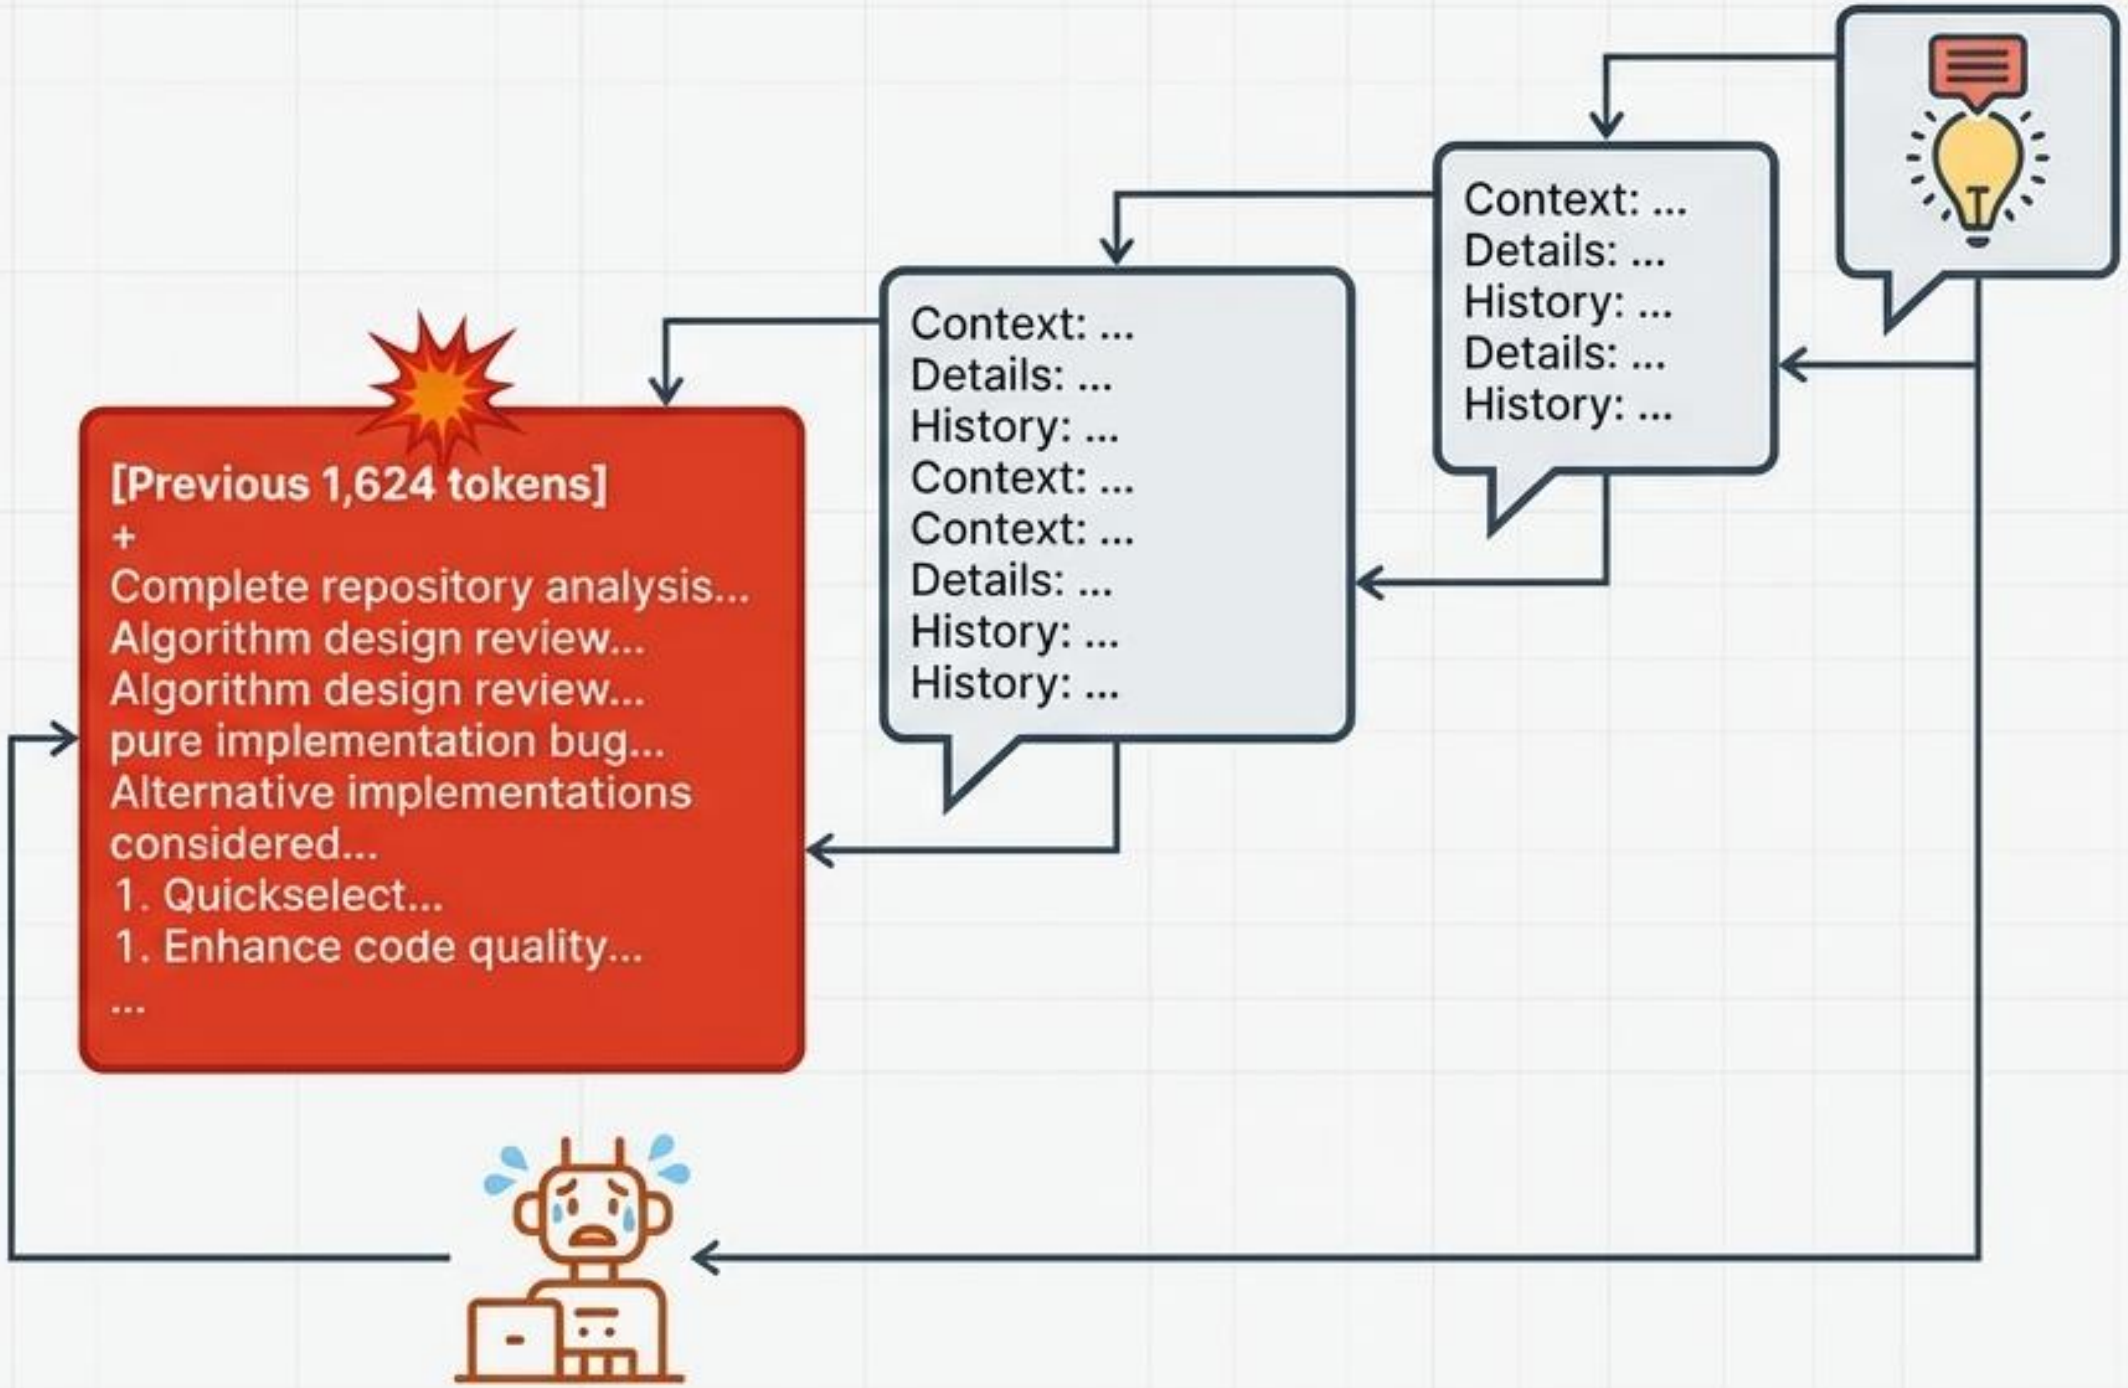

# Pathology 2: The Vanishing Textual Gradient.

**Symptom:** Diluted guidance and decaying specificity.

**Mechanism:** To manage context limits, compression and summarization strategies are applied at each hop.

**Result:** Critical anchors (file paths, line numbers) are stripped away, leaving non-actionable advice.

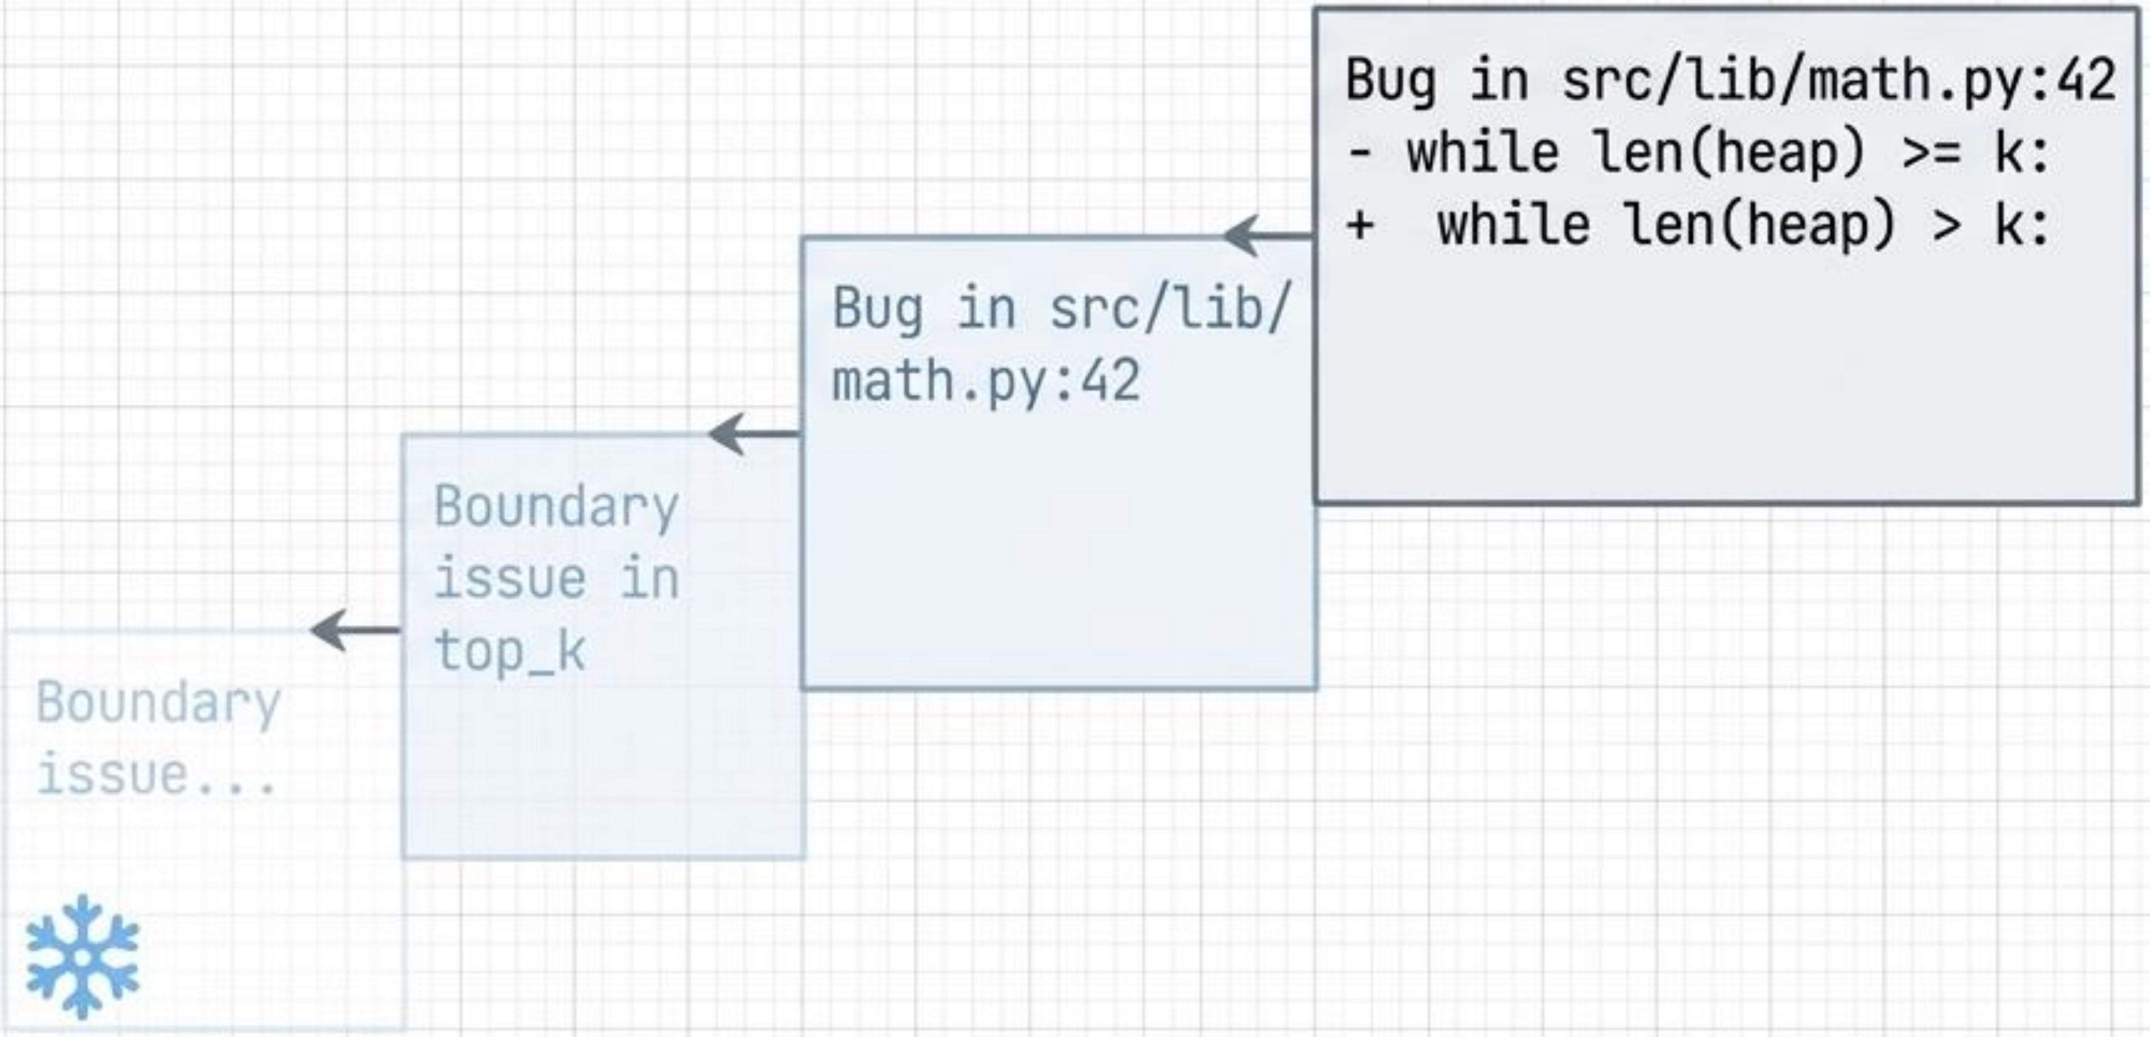


Step 1: "Bug in src/lib/math.py:42"

Step 2: "Boundary issue in top_k"

Final Step: "Consider refining comparison logic; edge cases need care."

# The Diagnostic Profile of Textual Optimization Failures.

| Failure Mode | Root Cause | End Result |
| --- | --- | --- |
| Exploding Textual Gradient 🔥 | Unbounded context accumulation to preserve execution traces. | Exponential token growth; "Lost in the middle" effect; Amplifies evaluation biases. |
| Vanishing Textual Gradient ❄️ | Over-compression to manage strict LLM context limits. | Strips exact anchors (lines/files); Yields non-actionable, generic advice. |

# The Paradigm Shift: Optimization must be localized before it is synchronized.

## Global Feedback Chains

Inherently unstable for language.

## Local Equilibrium

The Inspiration: `Energy-Based Models` use ‘`Equilibrium Propagation`’ to bypass the biological implausibility of `backprop`. We adapt this principle to bypass the contextual impossibility of textual backprop.

# Textual Equilibrium Propagation (TEP)

A local learning framework for deep agentic workflows.

**Phase 1: Free Phase**

Local LLM critics iteratively refine prompts at each node until reaching an 'equilibrium' independently.

---

**Phase 2: Nudged Phase**

Proximal prompt edits with bounded modification intensity applied via forward signaling.

**Advantage:** Completely avoids the computational burden and signal degradation of global backward chains.

# Mechanism 1: Reaching local equilibrium (The Free Phase)

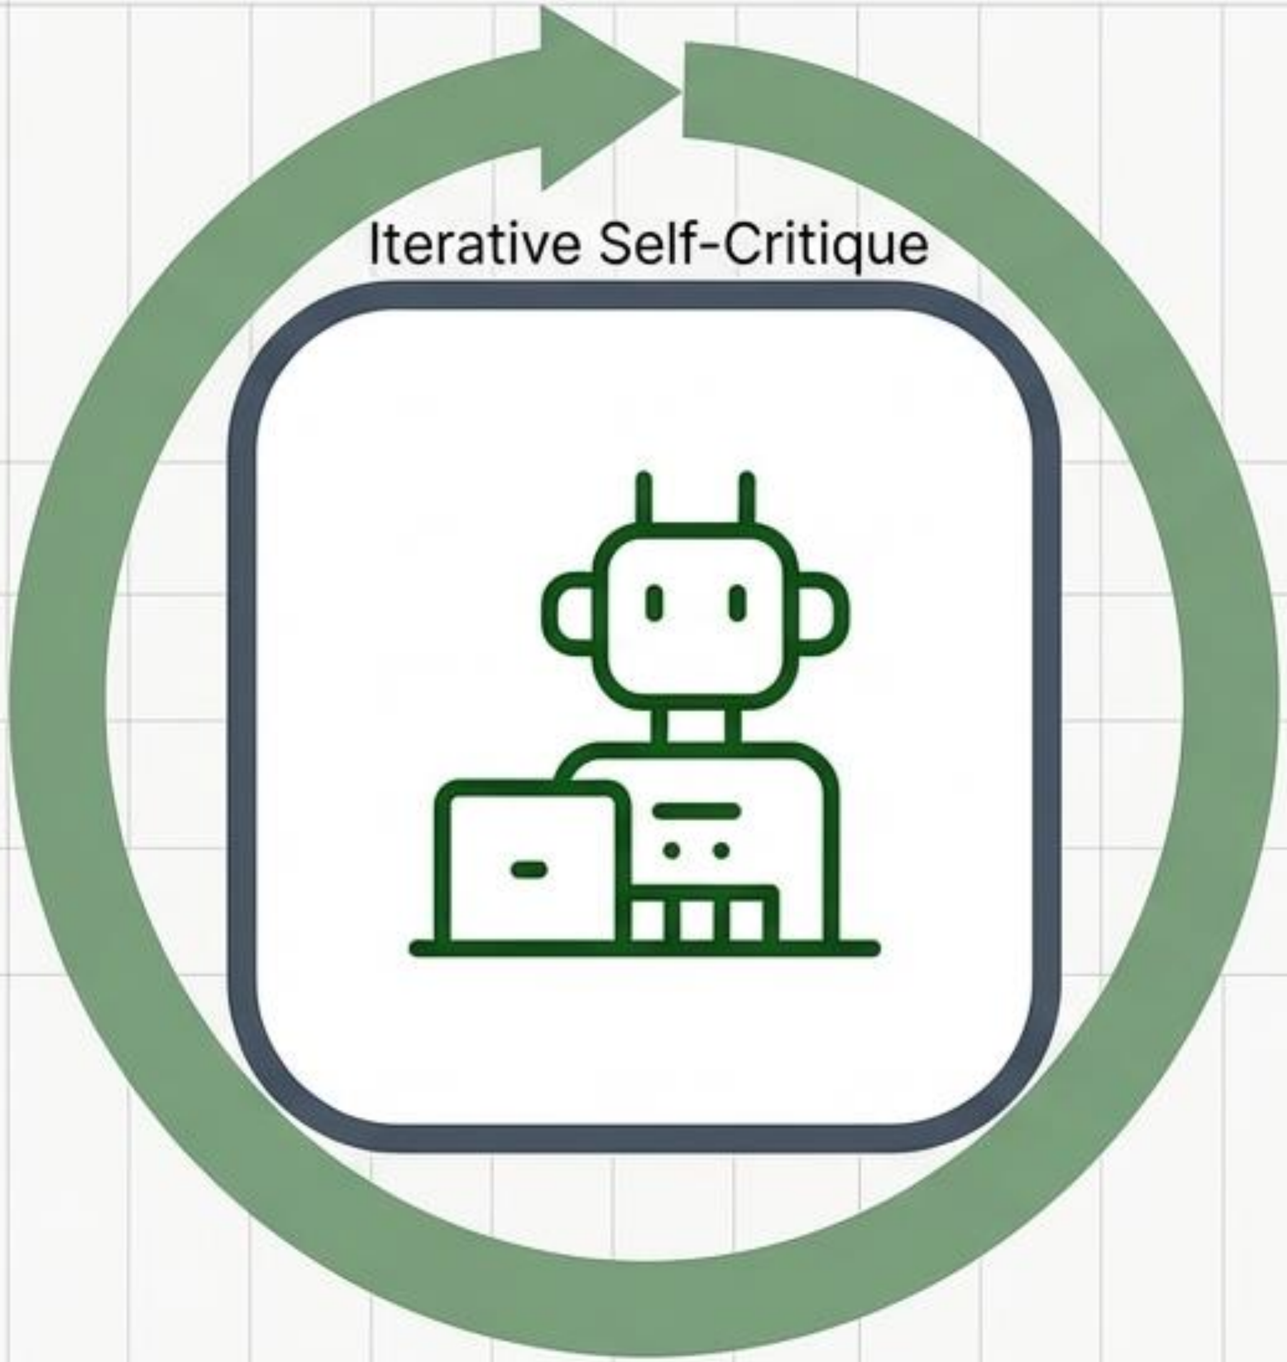



1. **Isolate the Node:** The agent is completely decoupled from the global backward chain.
2. **Iterative Critique:** Local LLM critics review the node's specific inputs and outputs independently.
3. **Achieve Stabilization:** Prompts are refined repeatedly until the local critic suggests no further edits.

# Mechanism 2: Bounded forward adaptation (The Nudged Phase)

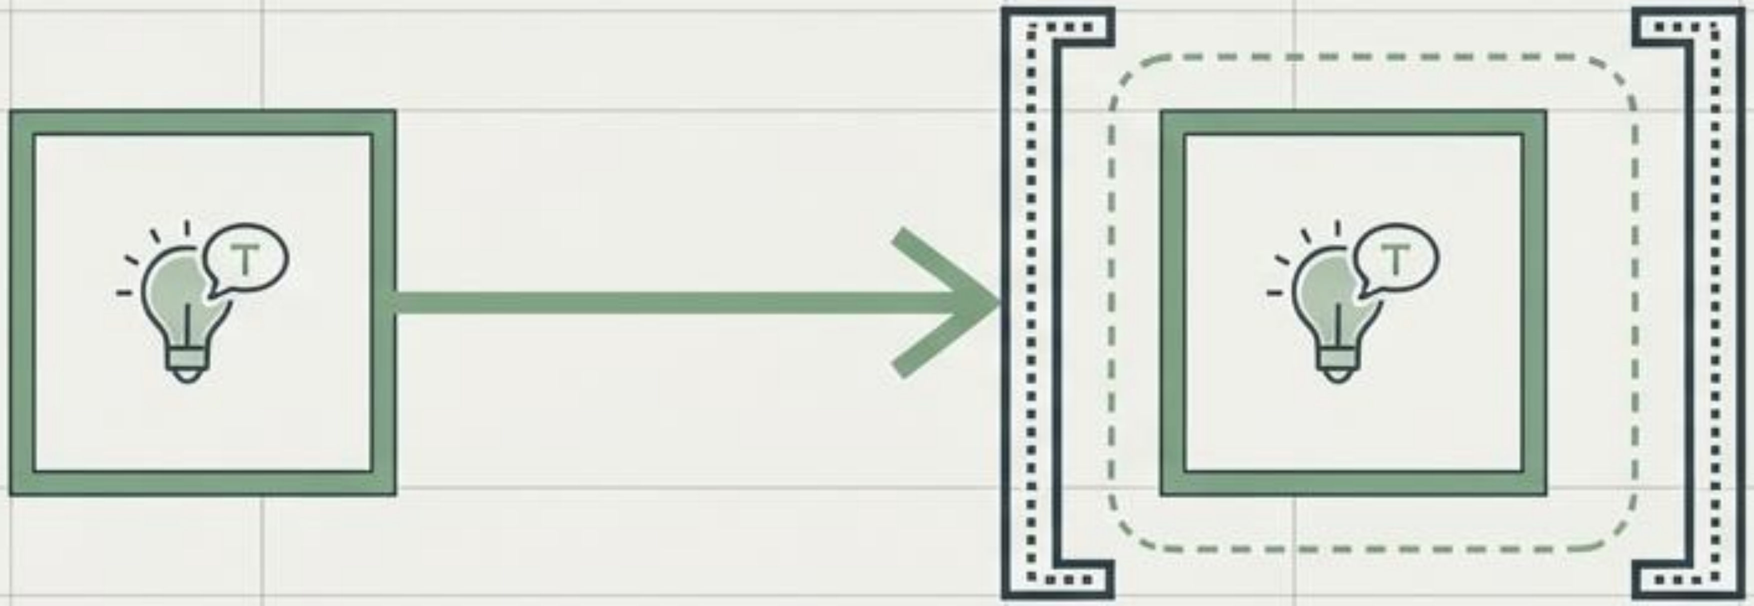

### Forward Signaling

Task-level objectives are propagated forward, never backward.

### Proximal Edits

The system applies gentle 'nudges' (minor prompt edits) toward the global target.

### Bounded Intensity

Modifications are strictly clamped inside the bounded fence to prevent chaotic rewriting and token explosion.

**Result:** Controlled adaptation toward global goals while preserving local node stability.

# The Optimization Paradigm Head-to-Head

| Dimension | TextGrad (Baseline) | TEP (Proposed) |
|---|---|---|
| Core Mechanism | Global Textual Backpropagation | `Local Equilibrium + Forward Nudge` |
| Update Direction | Long backward feedback chains | **Forward signaling** |
| Depth Scalability | Degrades exponentially (Explodes/Vanishes) | Maintains **stable optimization signal** |
| Computation Burden | High (requires re-evaluating full chain) | Low (bounded local edits) |
| Component Compatibility | Requires specialized tracking | Preserves **practicality** of black-box LLMs |

# The future of Compound AI requires stable, local learning.

### 1. Depth Demands Stability

Single-agent LLMs are giving way to deep, long-horizon multi-agent workflows (WebArena, ToolBench).

### 2. Global Backprop is a Dead End

`TextGrad`-style global chains will inevitably shatter under the weight of exploding context and vanishing specificity.

### 3. TEP is the Inevitable Architecture

By securing local equilibrium first, TEP improves accuracy and efficiency, with gains that actually grow as workflow depth increases.