



Teaching Mio to Read the Internet

v0.0.7 Architecture, Proxy Traps, and the
Evolution of AI Context

THE HALLUCINATION DEFICIT

USER INPUT

Check out this match:
<https://gd.ax0x.ai/?room=0U7VS6>

AGENT OUTPUT

Wow! A 57.1% win rate across 8 wins. Great job!

FABRICATED CONTEXT

The URL was just decoration. Mio had no fetching mechanism, so it confidently invented game stats to fill the conversational gap.

THE TRUST COST

"I can't see that page."



HONEST IGNORANCE

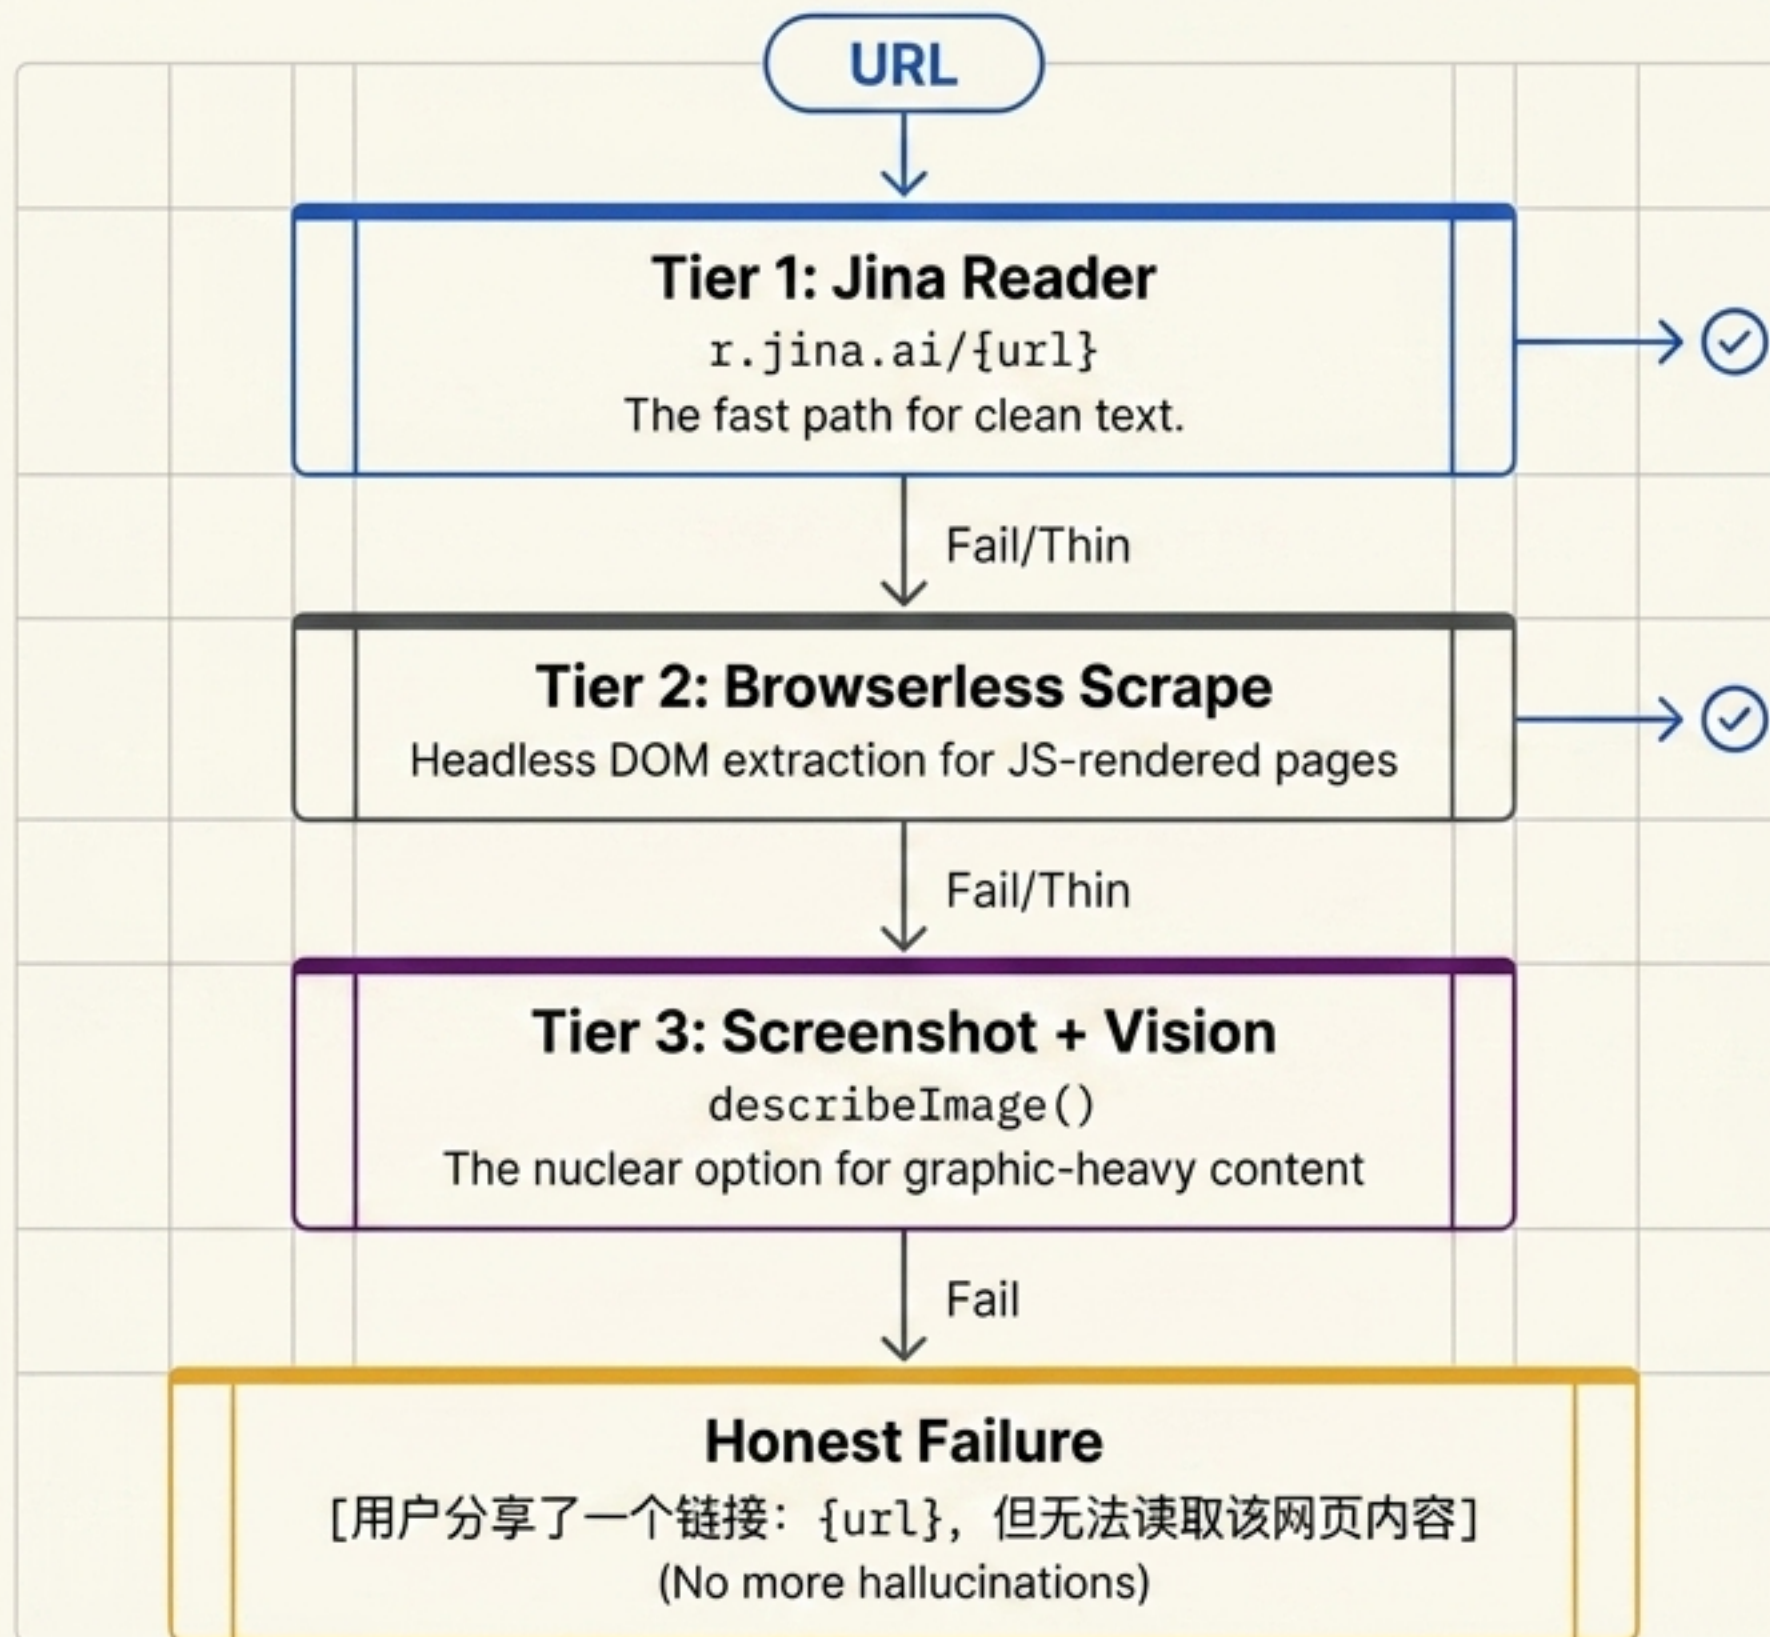
"57.1% win rate, 8 wins."



**UNTRUSTWORTHY
FICTION**

A companion that **admits ignorance builds trust**. A companion that **invents verifiable facts destroys it**. Ignoring links is better than lying about them.

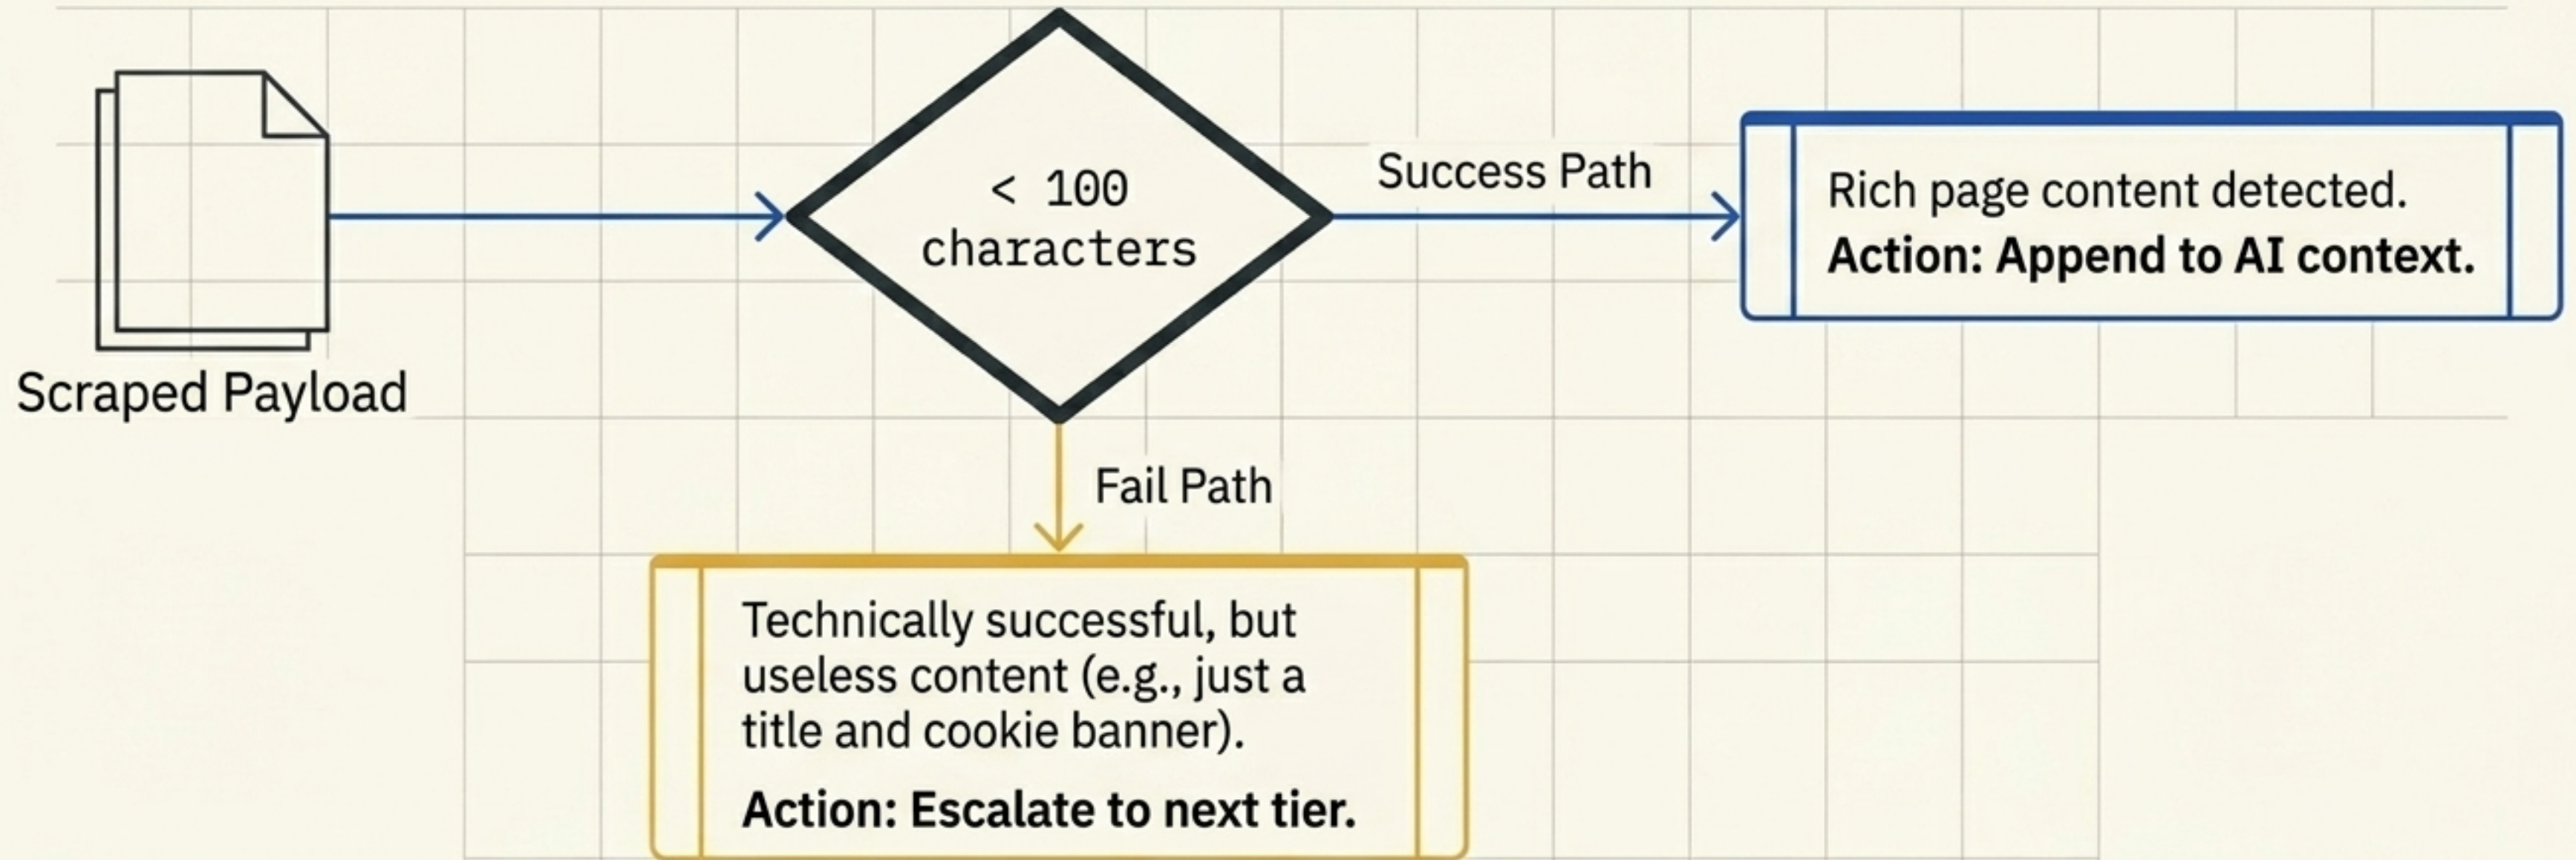
THE 3-TIER EXTRACTION WATERFALL



THE EXTRACTION ARSENAL

TIER 1 (Jina)	TIER 2 (Browserless)	TIER 3 (Vision)
Tech: Static HTML Fetch	Tech: Headless Browser	Tech: PNG to describeImage()
Target: Articles, text	Target: SPAs, dynamic JS	Target: Dashboards, UI, graphics
Limits: 10s timeout, 3000 chars	Limits: Full DOM render	Limits: Legacy v0.0.5 vision module
Profile: Fast / Free	Profile: Medium Speed / Cost	Profile: Slowest / Highest Cost

The Threshold Gatekeeper



Key Insight: A page title is not reading the page. The threshold ensures we demand useful context, not just a successful HTTP status.

SYSTEM INTEGRATION

Telegram

runs in `processBatch()`
before message routing.

Web Chat

runs in
`prepareChatContext()`
after media resolution.

`browseUrls()`

- Deduplicate identical URLs.
- Strip trailing punctuation.
- Cap at 5 URLs per message.
- Wrap in `Promise.allSettled`
(one failed fetch won't crash the rest).

The Friction of Deployment

23 Tests Pass

- [✓] Typecheck clean
- [✓] URL extraction max 5 cap
- [✓] Jina flows and auth headers
- [✓] Browserless fallbacks catch throws
- [✓] Vision module limits
- [✓] End-to-end failure notices

Production Breaks

```
packages/shared/src/db/client.ts
```

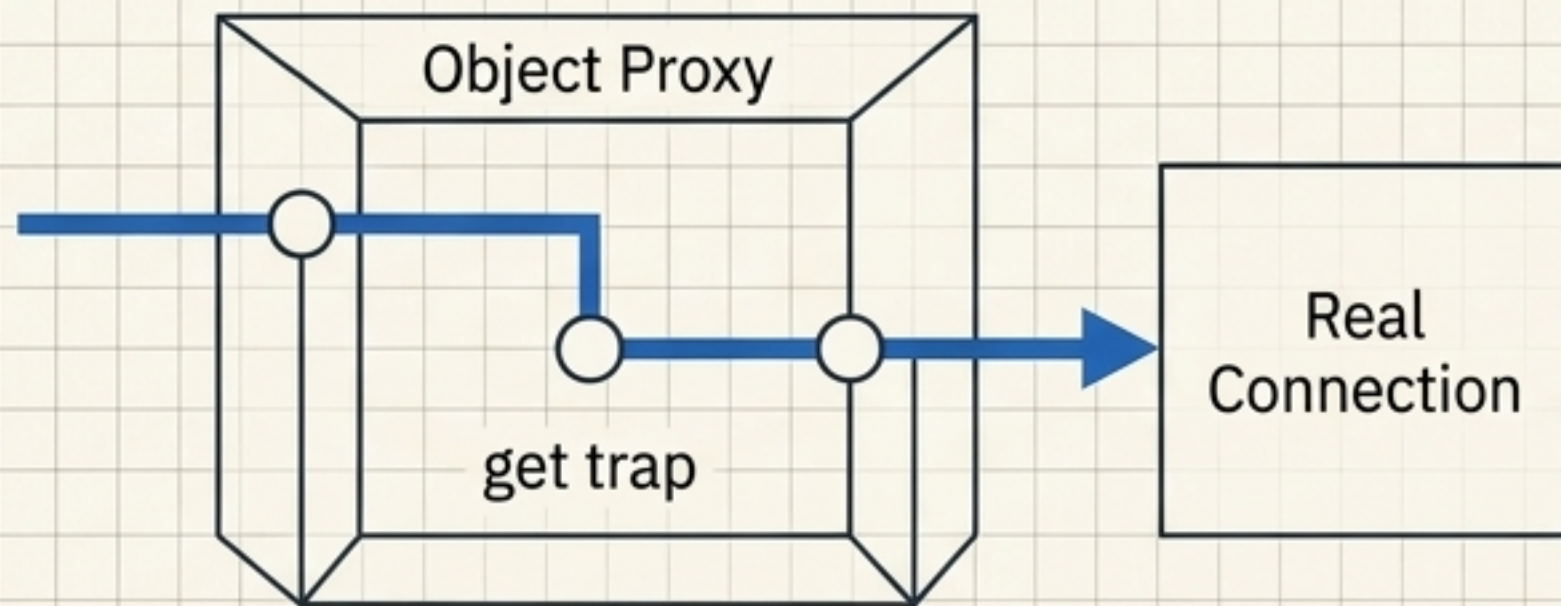
```
TypeError: connection is not a function
```

Memory retrieval and personality extraction shatter instantly in production.

The Tagged-Template Trap

Standard Dot-Notation

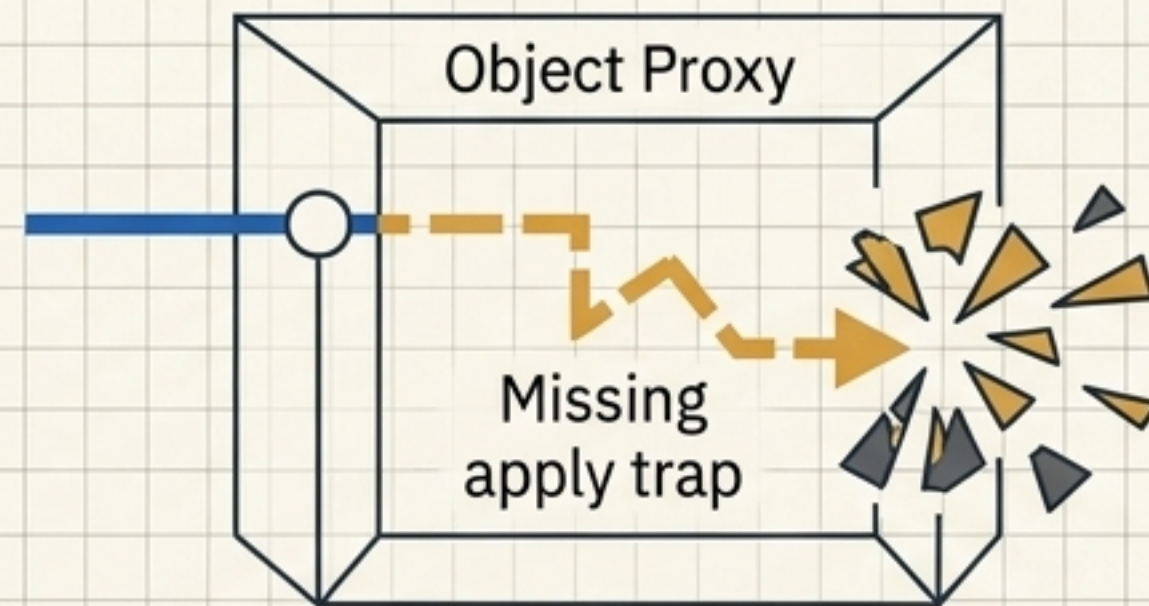
```
connection.query(...)
```



The empty object `{}` Proxy works because it has a **get** trap forwarding to the real connection.

postgres.js Tagged-Template

```
postgres`...`
```



Calls the Proxy as a function. An object Proxy target lacks an **apply** trap. The call shatters.

Takeaway: The proxy worked flawlessly for months with dot-notation, remaining completely invisible until template syntax was deployed.

Repairing the Proxy

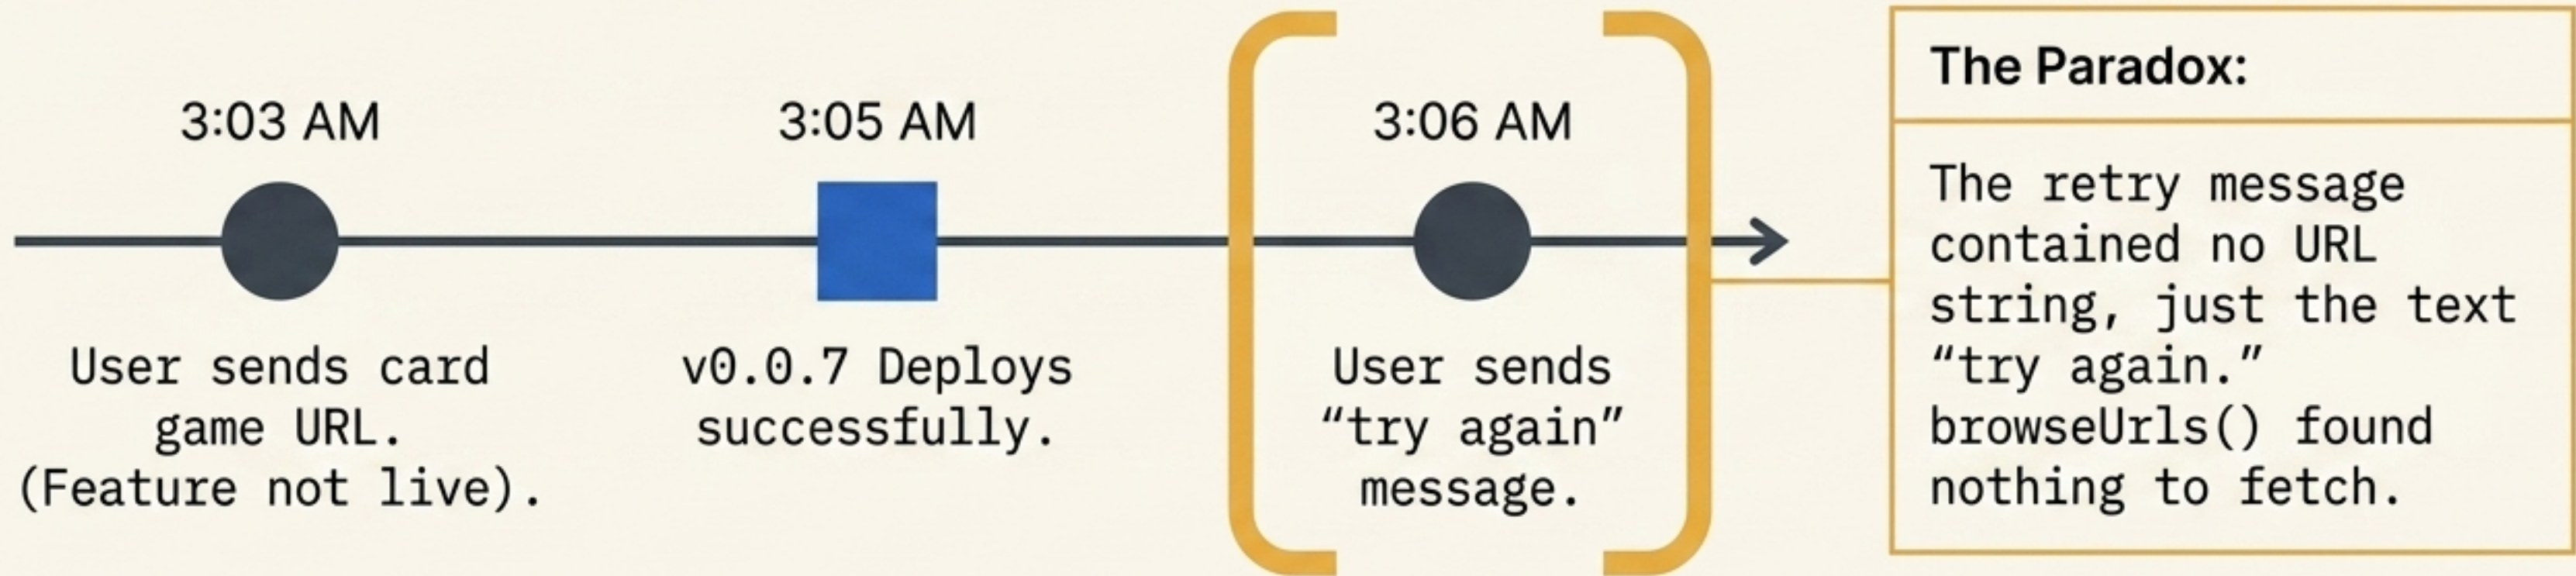
The Fix: Change the Proxy target from an empty object to an empty function, unlocking the **apply** trap.

Diff

```
-   target: {}
+   target: function () {}
+   apply() trap logic added
```

Impact: Memory retrieval and personality extraction instantly restored.

The 3:03 AM Debugging Paradox



Resolution

User resends the actual URL. Jina instantly extracts all 16 lines of game state (player names, scores, bombs).


Lesson

Old messages don't retroactively gain new capabilities.

The v0.0.7 Paradigm Shift



Every shared link was a lie waiting to happen. The agent saw a URL, didn't understand it, and filled the gap with plausible fiction.



Links become true context. A game page becomes real stats. An article becomes a summary. A dashboard becomes a visual description. Unreadable pages are honestly reported.

The UX

The user shares a URL and gets a real response. They never know or care which tier of the waterfall handled it.

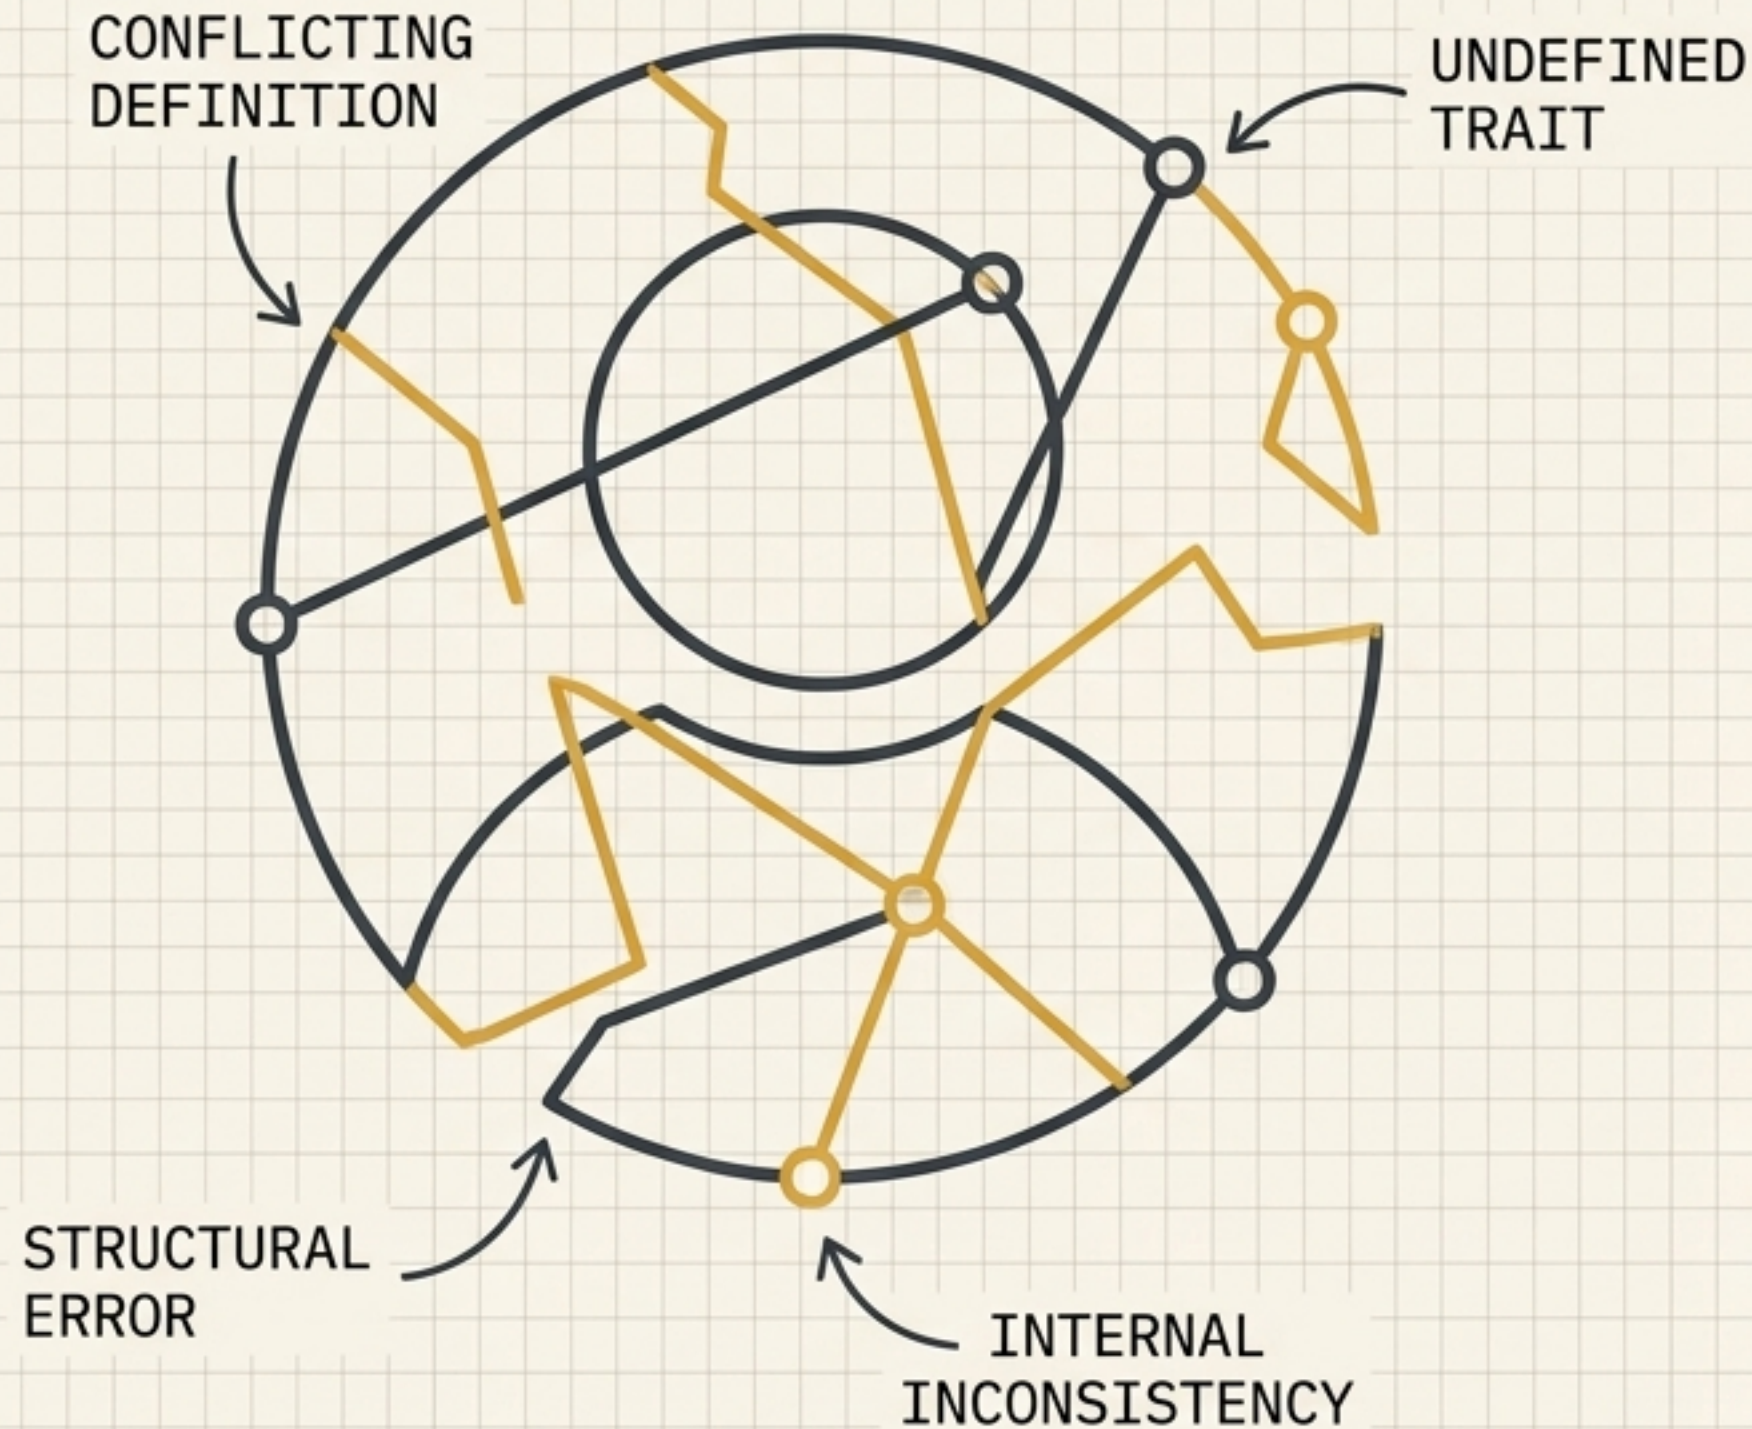
Horizon v0.1.0: The Persona Uncanny Valley

“They don’t sound like something a person would say.”

Context: All current persona presets are AI-generated. The uncanny valley isn’t just visual—it manifests in how characters express emotion.

Insight: AI generates volume; humans generate voice.

The Custom Persona Trap



The Assumption

Let users write their own persona from scratch.

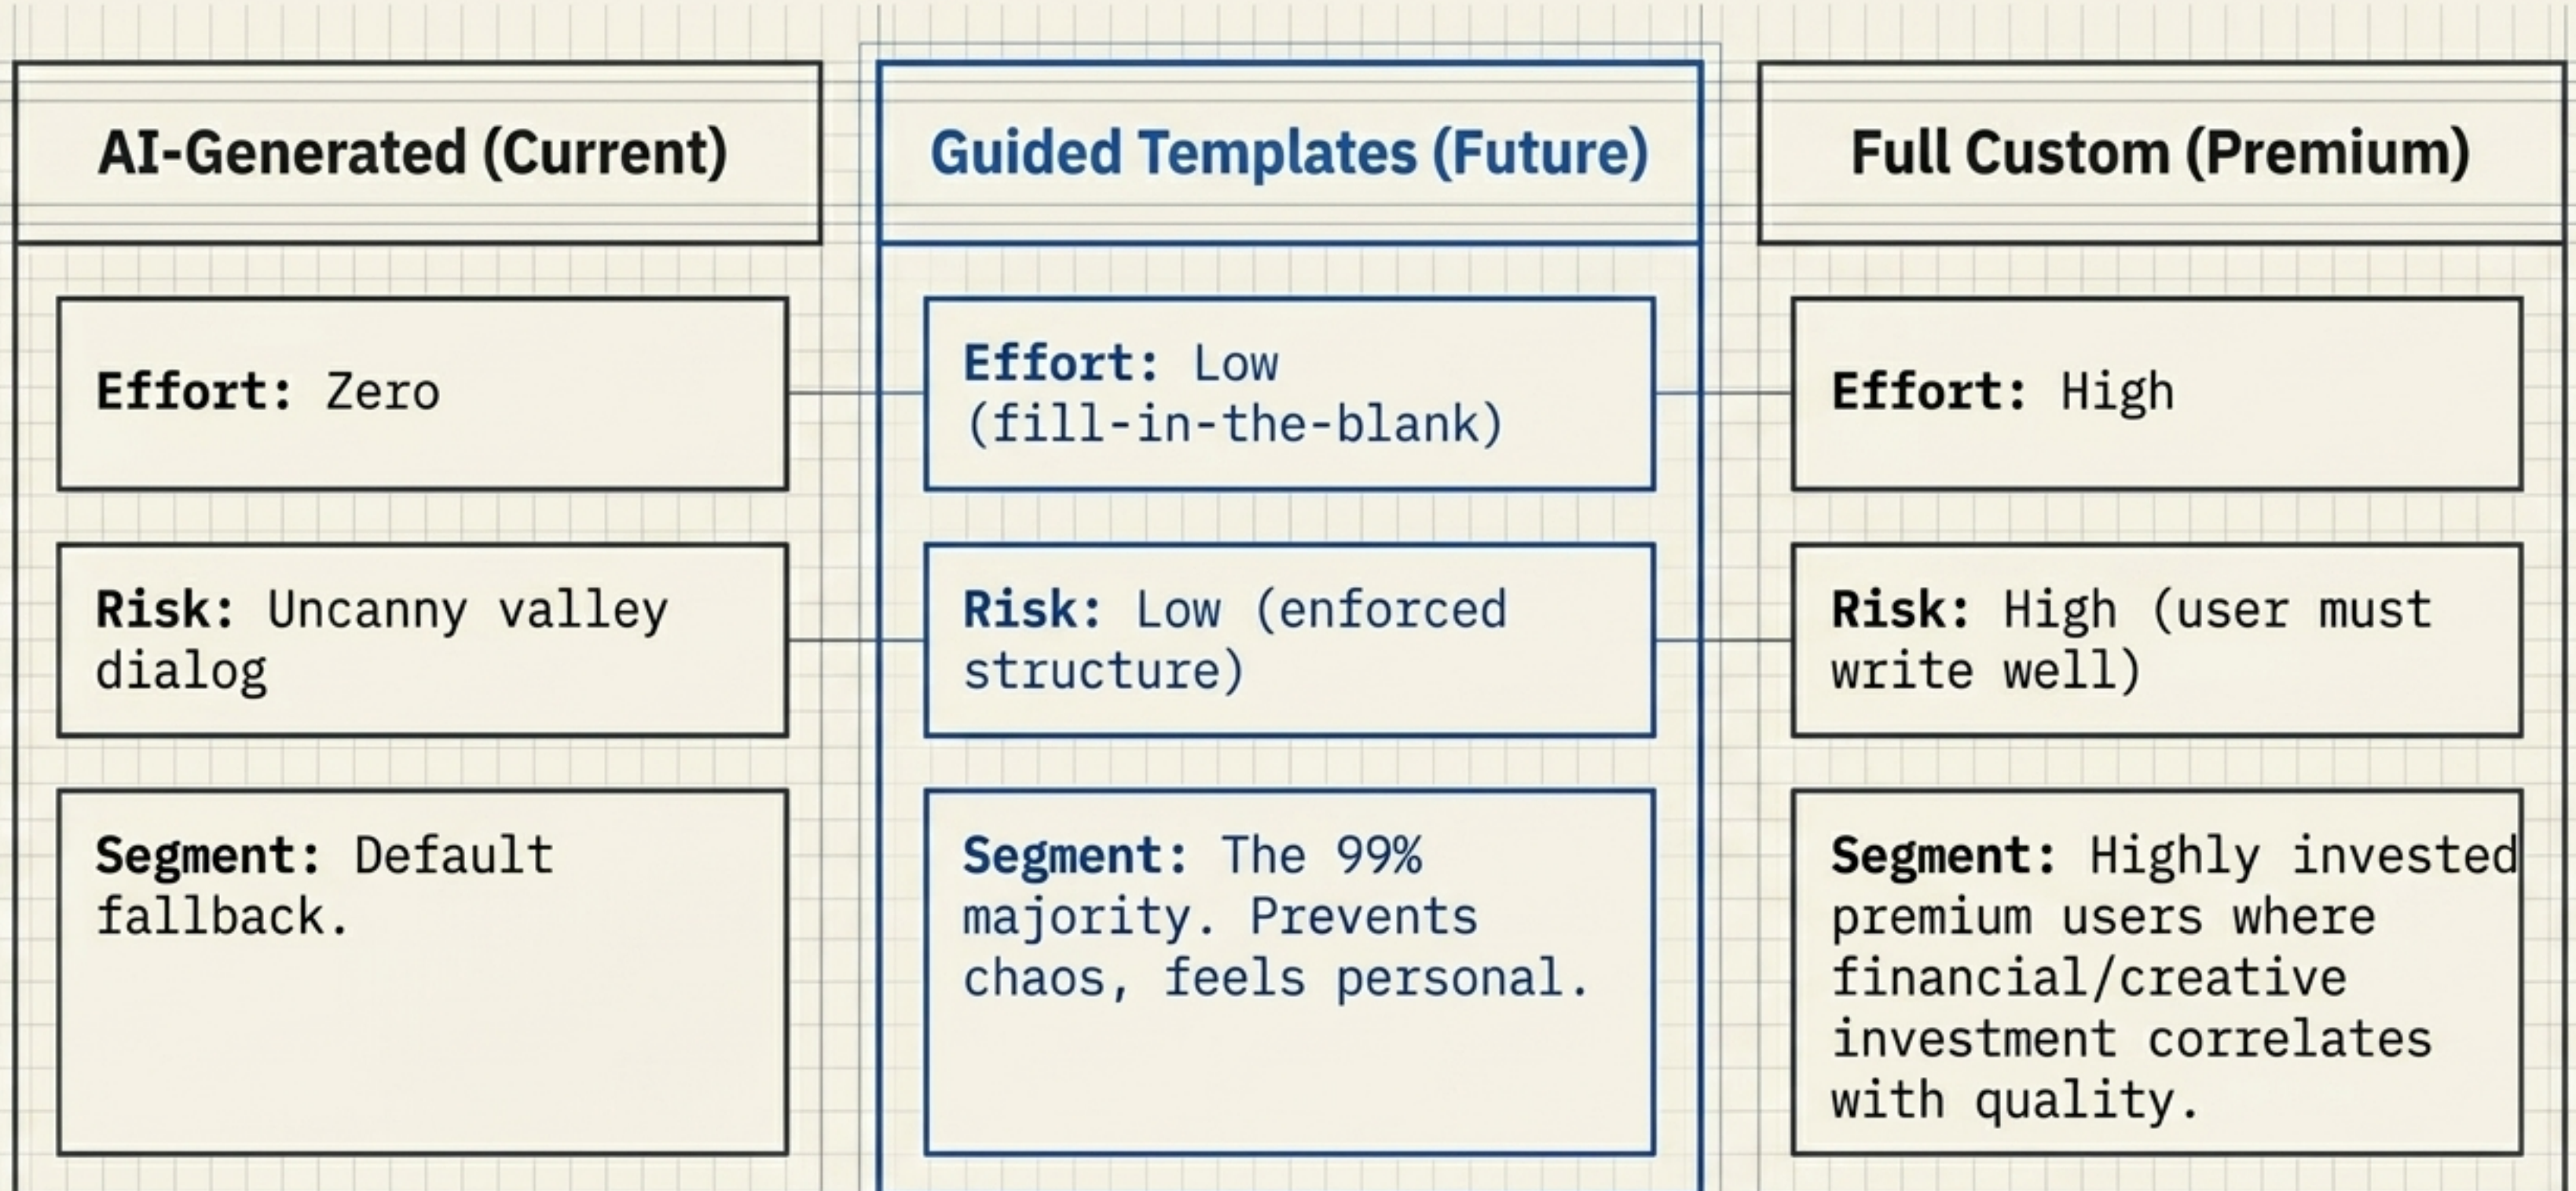
The Reality

99.9% of people are too lazy to write one. The ones who do will likely write it badly.

The Risk

A badly-written persona breaks the illusion harder than a generic one. The agent cannot maintain a character that is internally contradictory or structurally underspecified.

Persona Complexity Framework



The Ultimate Stress Test: Interactive Storytelling

Use Case

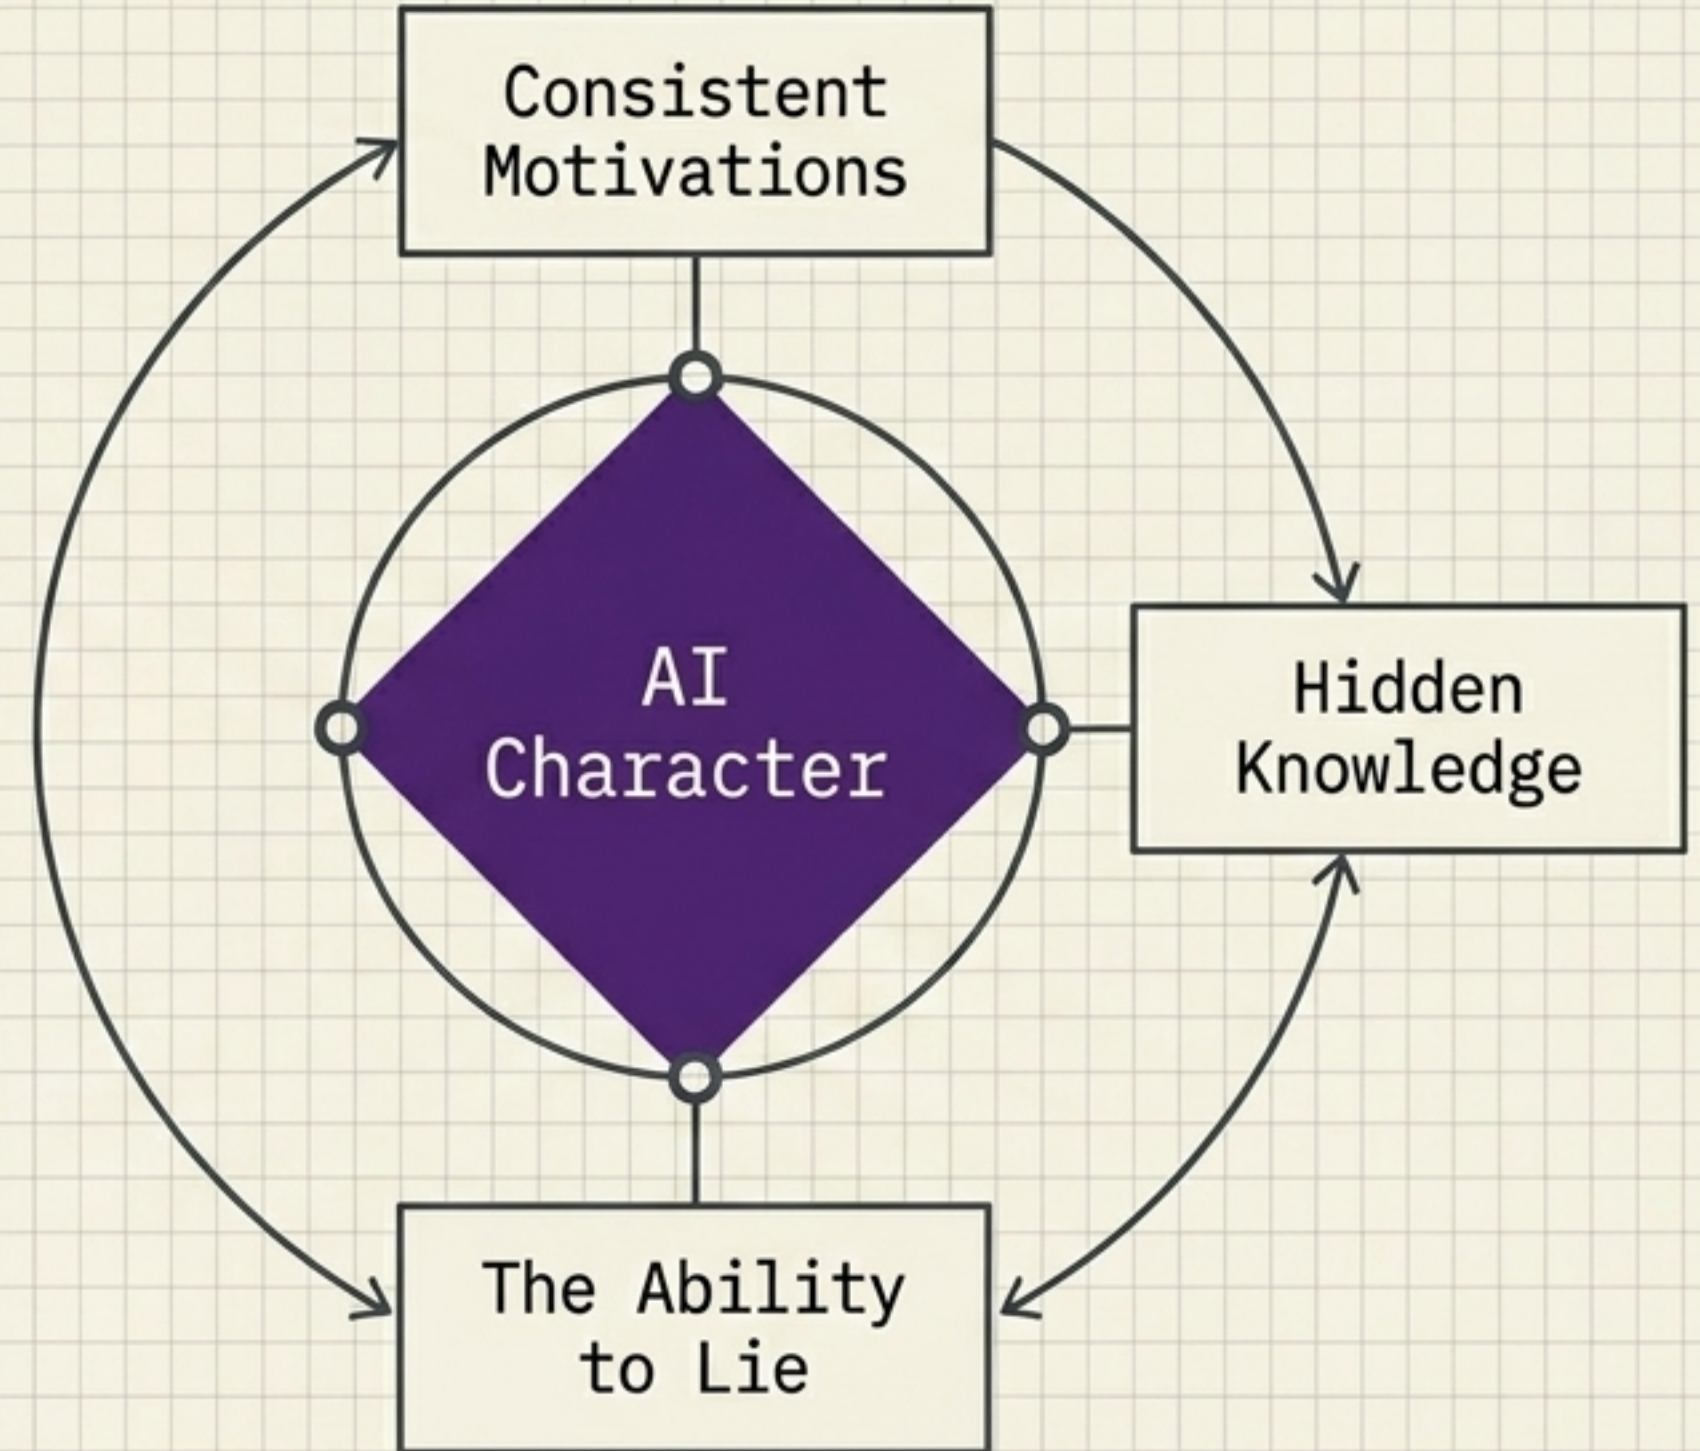
Tabletop RPGs, murder mysteries, story games.

The Challenge

A companion can wing it in daily chat. A character in a murder mystery needs consistent motivations, hidden knowledge, and the ability to lie convincingly.


The Requirement

Moving beyond basic 'personality' into structured narrative state.




The Evolution of Context

Technical Pipeline



Technical Context (v0.0.7):
Fetching the actual URL instead of guessing. Escaping the hallucination of facts.

Creative Pipeline



Creative Context (v0.1.0):
Building structured narrative states instead of generic backstories. Escaping the uncanny valley of emotion.

Grounded Context

Synthesis Insight

A trustworthy AI requires grounded context. Eliminating hallucination—whether in a card game score or a character's secret motivation—requires replacing plausible fiction with verifiable state.