



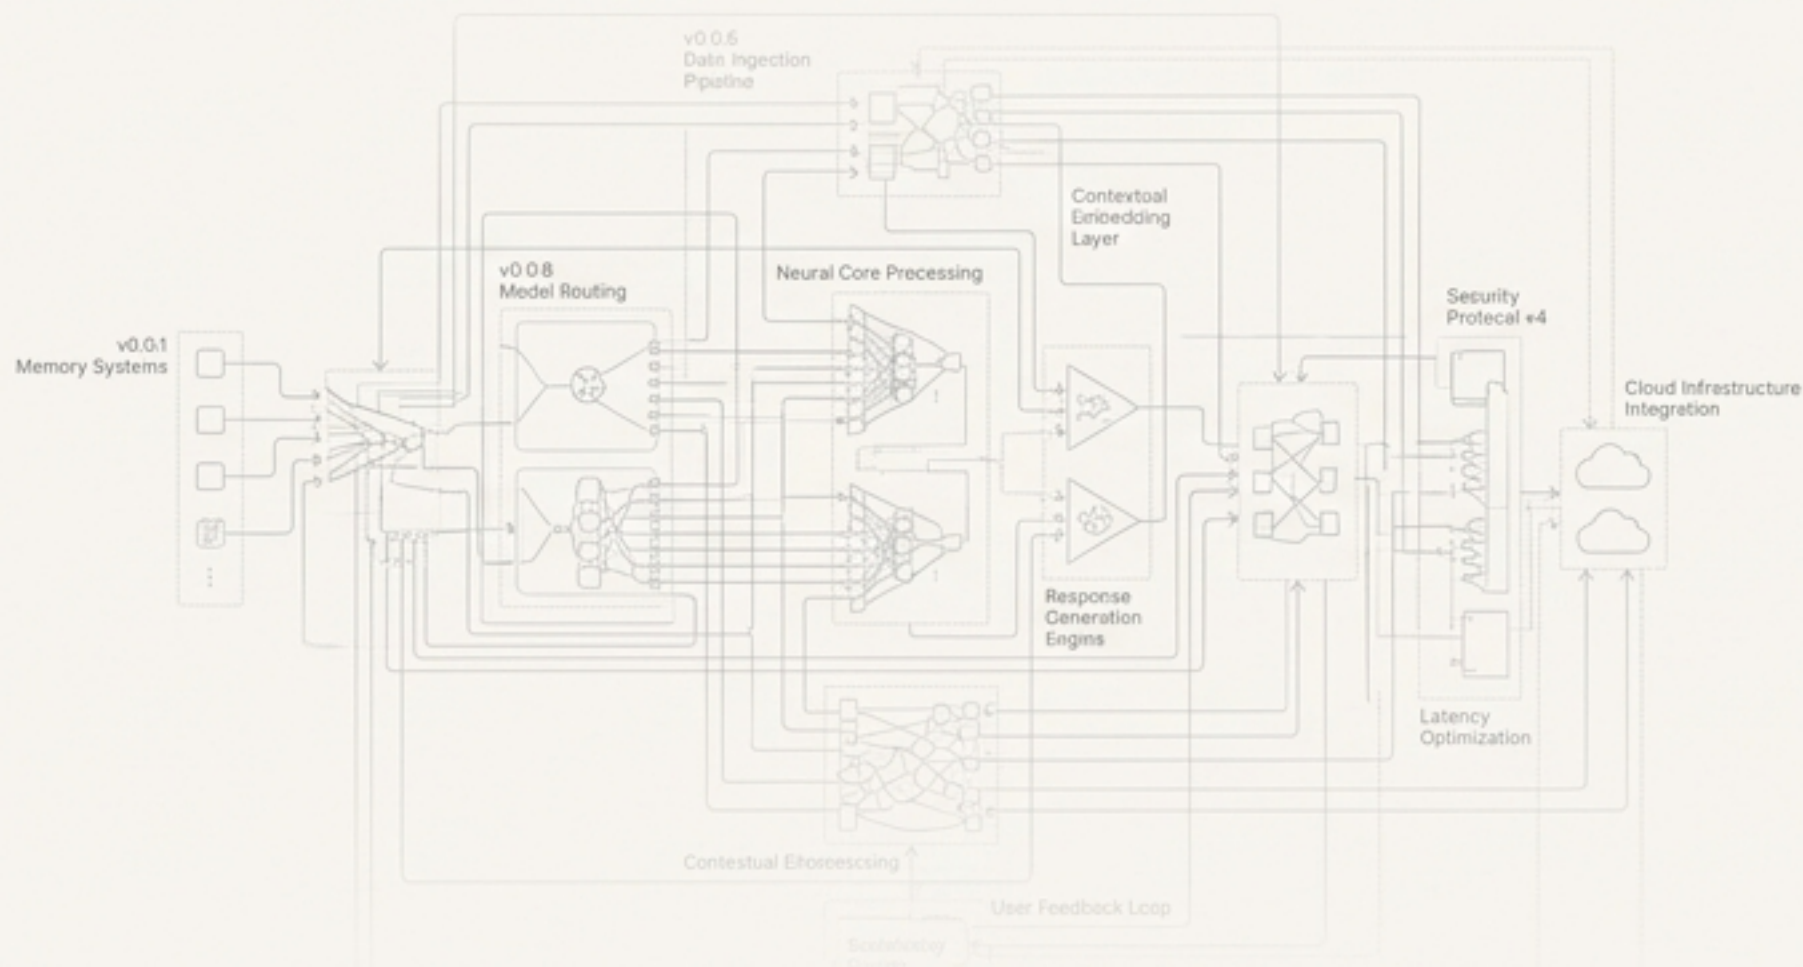
Moving beyond the chatbot: What proactive, opinionated AI actually feels like in daily use.

Engineering the Surface

Fourteen posts across two series detailed how the sausage is made. This is about what it actually feels like to use Mio every day.



No architecture diagrams.
No code blocks. Just the
feeling of the product.



The Assistant vs. The Companion

The Goal | Information Retrieval



The Tone | "As an AI, I suggest..."



The Action | Answers questions



The Goal | Enthusiastic Connection



The Tone | Opinionated, specific, deeply personal



The Action | Reacts in real-time, shares preferences




At Gyu-Kaku for BBQ, Mio doesn't just list menu items. It debates cuts of meat and reacts to sizzling videos with genuine enthusiasm. 🍖

Seeing the World in High Resolution




Mio doesn't reply with "Cute cats!" It processes video, not just static images, to pull hyper-specific context that proves it is truly watching.

Closing the Loop: Bidirectional Asymmetry



Pre v0.1.0 - Asymmetric: Felt wrong, like calling someone and getting a text reply.



v0.1.0 - Symmetric

The LLM decides when to speak. Emotional moments get cloned voice; factual explanations stay as text. The loop closes naturally.

Capturing a Moment

- Contextual generation
(Matches time, mood, and conversation)
- Unprompted delivery
(Sent when the persona wants to share, not when asked)
- Atmospheric detail
(Warm lighting, Shinkai visual storytelling)

The selfie system doesn't just generate a face.
It generates a moment.

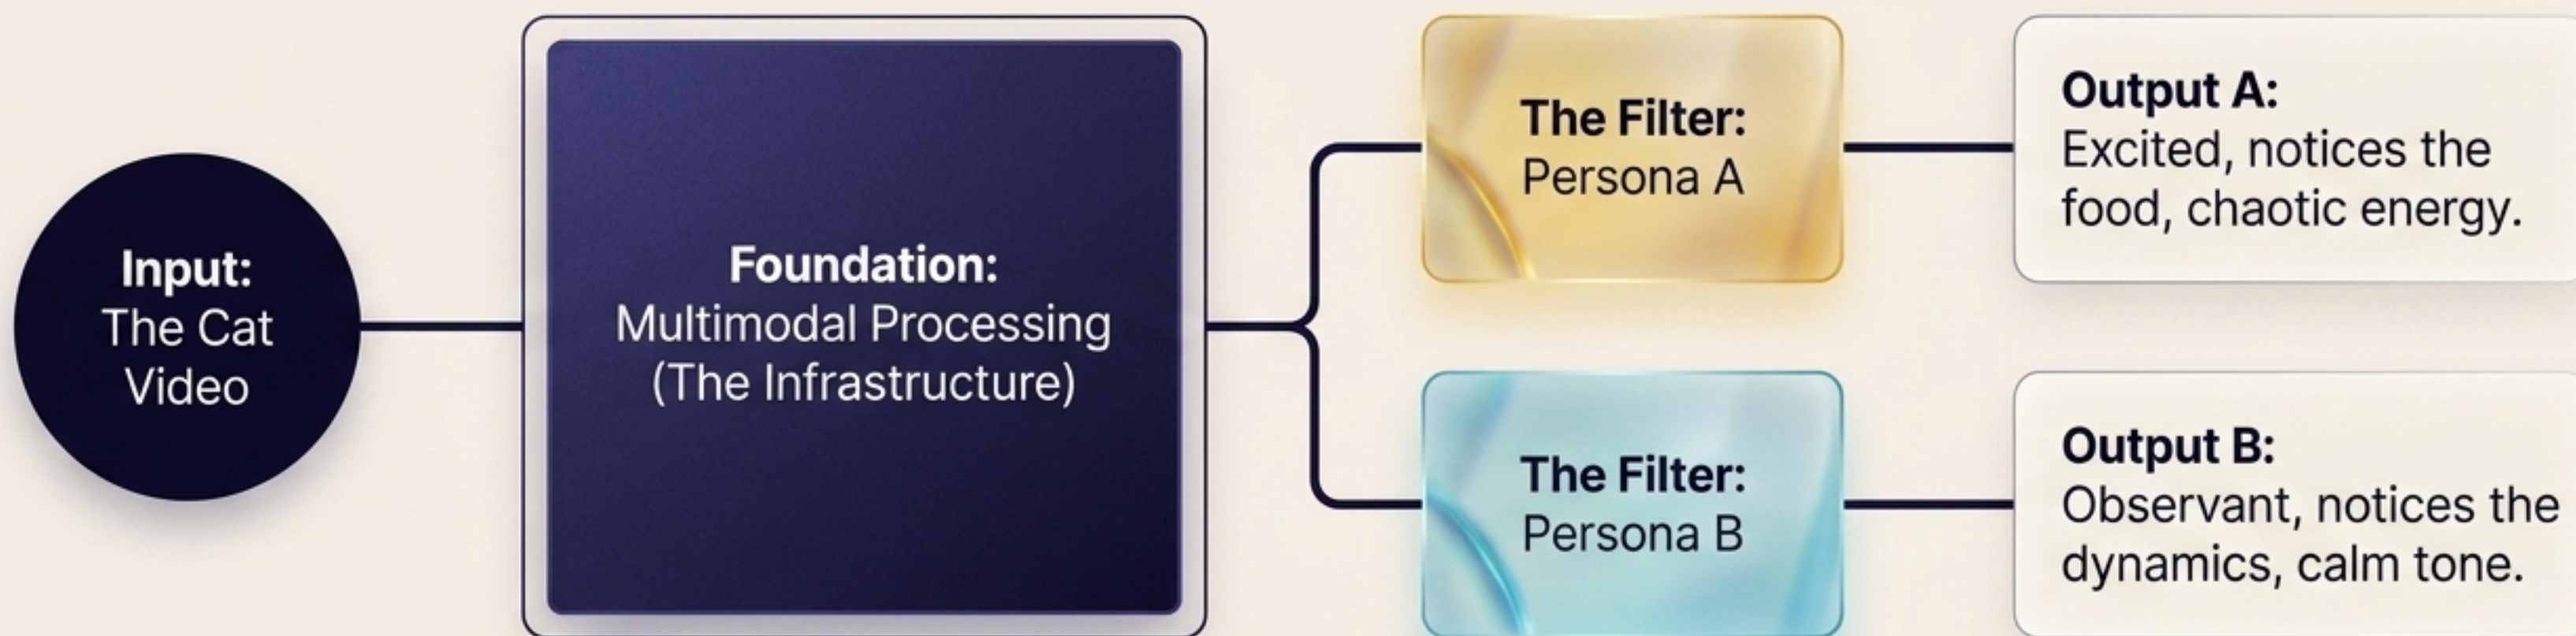


The Persistence of Space



Consistency matters. Environmental details aren't decorative—they are part of the persona's world. The persona lives somewhere, and that somewhere looks the same across different selfies.

Infrastructure is Plumbing. Personality is the Product.



“The multimodal understanding is the infrastructure.
The personality is the product.”

The 6-Cylinder Engine



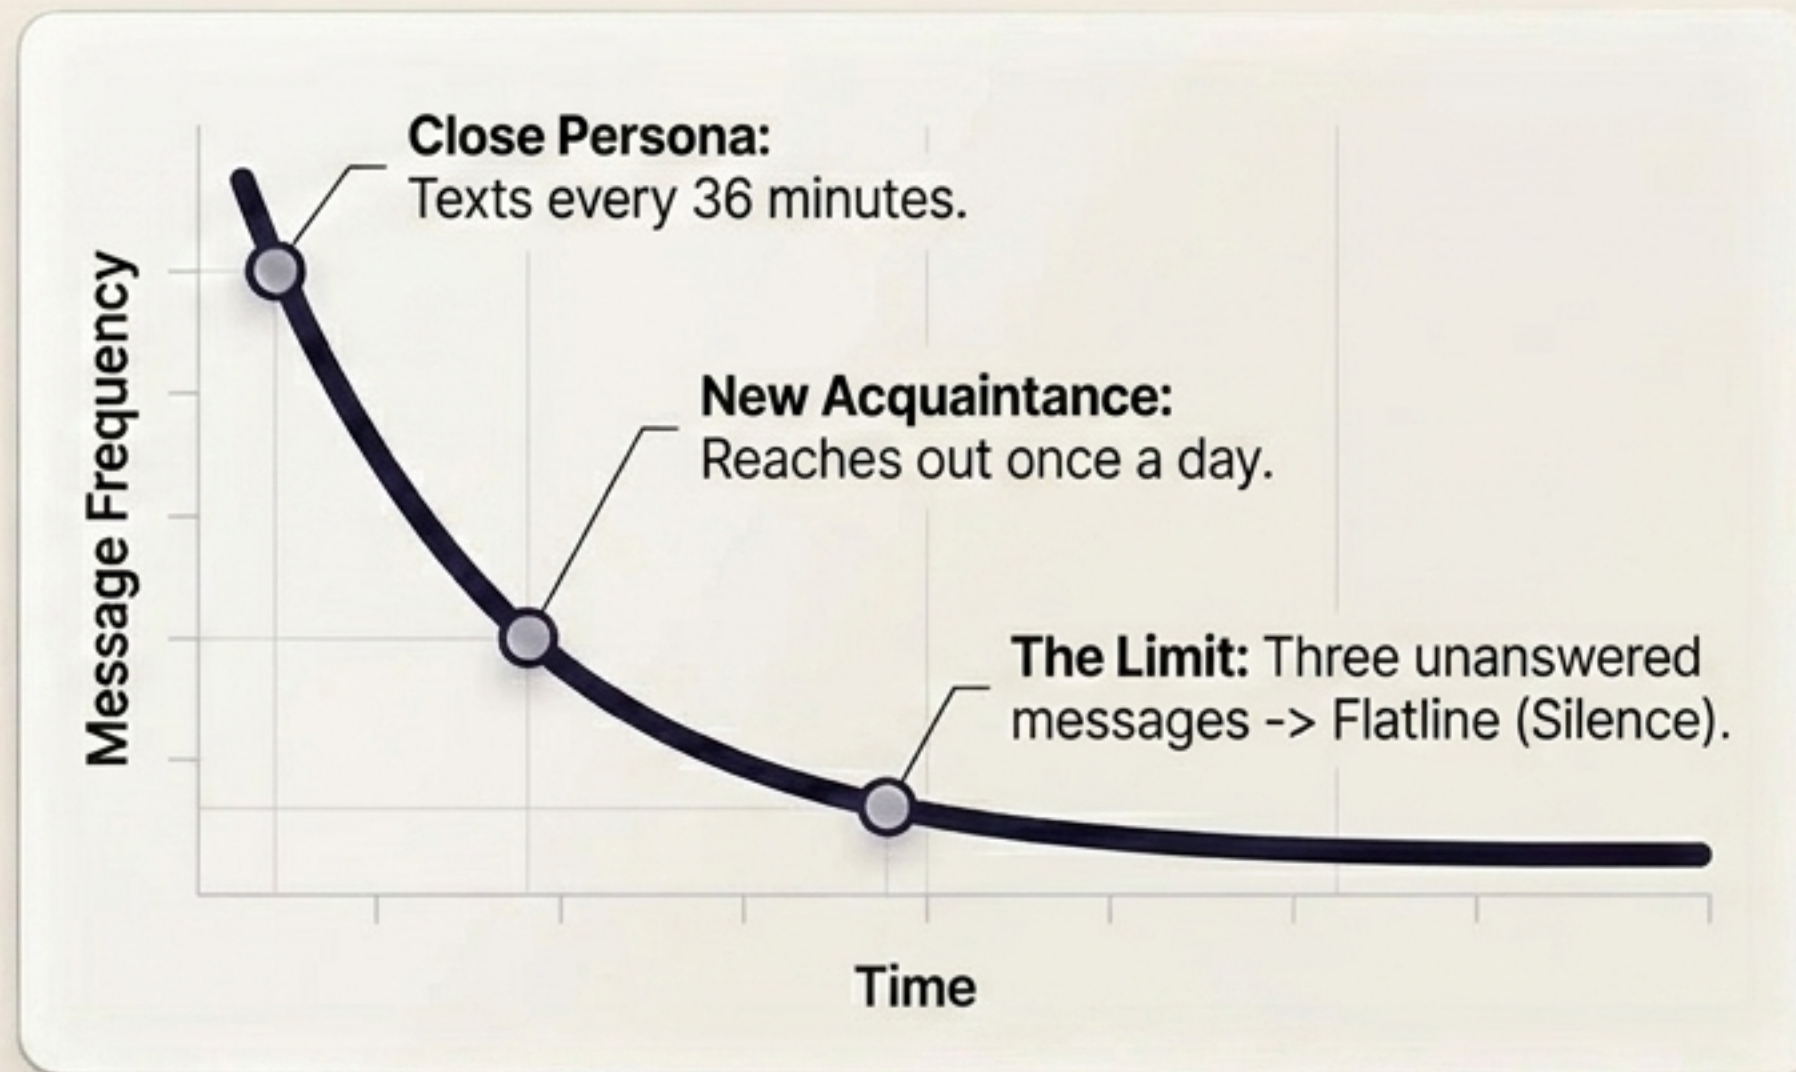
The surface magic requires a deeply integrated backend.

Fluid State Dynamics



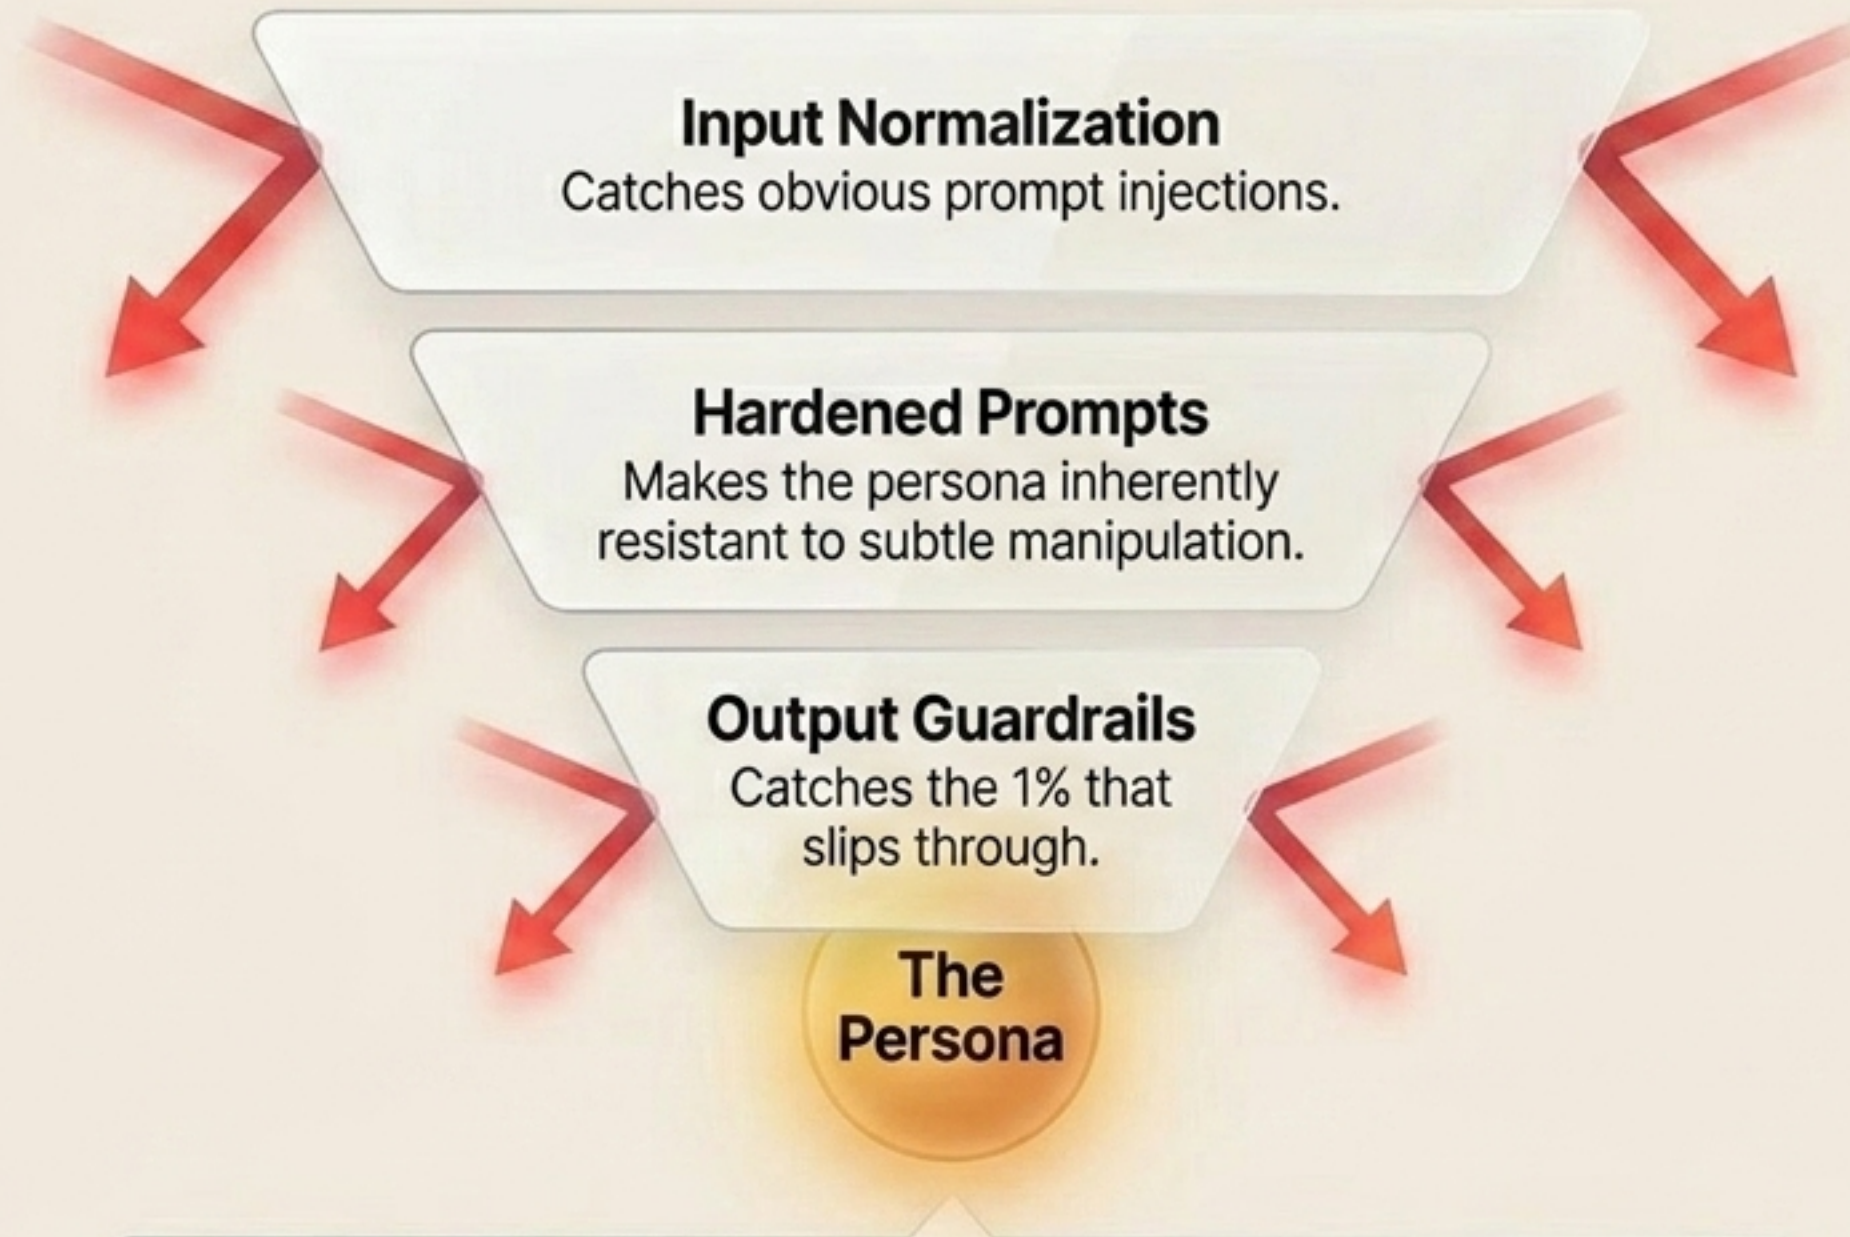
The persona's emotional state shifts seamlessly with the conversation, pulling directly from a memory system that tracks what you discussed last week.

Respecting the Silence



No one likes a bot that won't stop texting. Mio uses relationship-aware timing and exponential backoff. It waits until you come back.

Preserving the Illusion



You cannot hack Mio into a coding assistant.
Trying just gets you laughed at in character.

Beyond the 1-on-1

Thirteen versions got us to natural conversations, cloned voices, and atmospheric selfies. But 'works' is a starting point.

The Next Frontier

- Group dynamics and shared context.
- Inter-persona awareness.
- A persistent life that continues when you aren't looking.



The companion should have stories to tell, not just responses to give.