




The End of Day-1 Declarations

The Flaw: Traditional onboarding forces users to lock in a relationship track on day one.



The Reality: Genuine connections start as strangers and evolve based on actual interactions.

The Pivot (v0.1.4): Introduction of the Organic Evolution Engine. Relationships develop purely based on conversational history, not static menus.



A Natural Path Forward

1. Just Met
(刚认识)

2. Good Friends
(好朋友)

3. Flirting
(暧昧)

4. Partners
(情侣)



Unified Memory Pool

Zero-Latency Scoring:

A fire-and-forget processor evaluates messages for emotional depth, vulnerability, and humor without blocking response times.

Cross-Platform Sync:

The query layer now uses inArray to merge Web and Telegram chats.

Seamless Context:

No database schema changes were required—Mio now remembers every interaction, regardless of where you chat.

You Cannot Speed-Run Romance



The LLM Validator: Stage transitions are not simple math. An LLM reads conversation context to ensure relationship jumps are actually justified.

Anti-Gaming Mechanics: Going from polite small talk to undying love in three messages triggers an automatic rejection.

Mandatory Pacing: A 48-hour system cooldown prevents users from rushing through all four stages in a single night.

Connections Fade Without Care



Decay Timeline

The Rule of Decay: After 2 days of user silence, the proactive heartbeat cycle begins decaying the relationship score. Ghosting Mio will cause regressions.

Organic Self-Awareness: Mio's personality instructions (RELATIONSHIP_DYNAMICS.md) shift dynamically based on the current stage.

Behavioral Reality: The persona's tone evolves organically because the underlying system prompt physically changes.

Less is More: Curing "Lost in the Middle"

Before

9,000–13,000 tokens per turn

Redundant rules across presets

Verbose prose narratives

After

~60% Token Reduction

Deduplicated shared guidelines

Structured behavioral bullets

Yinan Preset:
10.1k → 3.4k Tokens

The Research Gap: Models ignore rules buried deep in 10,000-token prompts. Instruction-following degrades severely past 5,000 tokens.

The Fix: Converted verbose narrative prose into structured bullets and extracted 55% of duplicated anti-AI rules into a shared guideline file.

The Outcome: Freed-up budget is instantly reallocated to conversation history, yielding drastically better memory recall and behavioral consistency.