

Mountain of Many Voices

Mountain of Many Voices is a research project proposing a new kind of machine — one structured by its memory. This time, the machine is a story.

zzyw (Zhenzhen Qi, Yang Wang)

Hi, we're zzyw. I'm Yang, this is Zhenzhen. We are an art and research duo. We work WITH and ON computing — as a medium for making, and as a subject to be studied critically. Our practice moves between art, teching, research, writing, and coding.

The project we're presenting today, Mountain of Many Voices, short for “MOMV(move)”, is a research proposing a new kind of machine — one whose memory conditions what it can do. Here, it takes the form of a web-based hypertext narrative.

Before Zhenzhen walks you through the story and software, I want to spend a few minutes on the questions it's working with.

Software is infrastructure, but it is a mistake to treat it just as a technical substrate. It is also the ground infrastructure for the social, the cultural, and the epistemological.



Its structures and paradigms are becoming the limits of how open, rich, and generative thought and culture can be.

Infrastructure is produced, and it persists through time. In this sense, infrastructure is a form of memory.

The starting point may sound very familiar, that Software is infrastructure, but it is a mistake to treat it just as a technical substrate. It is also the ground infrastructure for the social, the cultural, and the epistemological.

Its structures and paradigms are becoming the limits of how open, rich, and generative thought and culture can be.

Infrastructure is produced, it persists through time. In this sense, infrastructure is a form of memory.



Memories are constantly reshaped by the present and the future, while also conditioning what the present and future can be.

If infrastructure is a form of memory, it's worth asking what kind of memory it currently is. Currently, memory in our technical systems is treated as inert: stored, and “queried” only when needed.

And, this points toward a new kind of machine: one that can afford us a freer and more open environment because it has a different relation with “memory.” Memory is never simply discrete; it cannot be separated out or categorized cleanly. It is not simply from the past, in the sense that memories are constantly reshaped by the present and the future, while also conditioning what the present and future can be. Memories exist in time, though **not** linearly.


What if systems had a more intimate and foundational relation with their technical memory?

extraction

id: 4471
content: "the river was higher
than I remembered"
timestamp: 2026-03-16T09:22:11Z
user: u_882

The utterance is preserved. The situation of its making is not part of the record.

herbarium sheet



The utterance is preserved alongside a record of the situation of its making. The format takes a position on what's worth carrying forward.

Inputs, outputs, history: their use is no longer plain retrieval from simple storage. They become part of the “context” — the condition for future creation and interpretation. AI enters here as connective tissue, interpreting this accumulated history in real time, through frames the writers and designers have authored.

Everything that has happened before becomes constitutive of the system *itself* — its functions, its features, its outcomes.

In practice, what comes next is neither a system entirely pre-completed by the authors or designers, as in traditional software—decision-tree logic, or, in a video game scenario, a branching plot—nor just “AI-generated content” that comes from nowhere and everywhere.

Memories, traces, inputs and outputs — they also carry intentions, orientations. When automation such an LLM enters, it must be conditioned not only by memory-as-storage but also by memory-as-orientation.

A way of seeing, thinking, and moving.

Data becomes heterogeneous and meaningful again, when its interpretation is held ... as a conscious act, against the history in which they were produced.

Mountain of Many Voices (MoMV) is a research-creation artifact developed alongside our inquiry into how memory might condition future operations of technical systems.

It is a Hypertext narrative system with a custom-made editor.



> One: Return



Thirty years old, Shane had come back. On the first day of 1996, he drifted like a ghost into the pale light of the valley as thin and translucent as

the plastic film on a sun-bleached pack of Marlboros. /
a cassette tape ribbon caught in a roadside briar. /
a moth's wing caught in the grill of a truck that hasn't moved since '92.

Just like the old times, the family kitchen smelled of stale coffee and cold ham. In the livingroom, the low-frequency hum of a television was left on

The story takes the shape of a pilgrimage where the player assumes the role of a thirty something protagonist returning to a home left behind by the rapid urbanization of the 1990s.

The journey begins as a grounded walk through streets thick with memory. Later on, as the player ascends the town's mysterious mountain, their individual recollections begin to entangle with other travelers, local wildlife, and spiritual entities.



Every player choice become the "contextual fabric" that dictates the technical generation of the future.


For example, in Chapter One, Return, The player completes the protagonist's arrival by selecting "A pack of Marlboros", "cassette tape ribbon" or a "moth's wing". This player input decides what the world is, and what is to be encountered.

By the time the player reaches **Spin-City Records**, the system expands on the narrative based on three things:

1. The specific metaphors the player has installed in Chapter One.
2. The specific version of the town's lore the player has uncovered thus far.
3. The system interprets this data through a specific interpretive framework—in this instance, Jungian psychology—to create a "wavy text" effect that marks the descent into the collective memory.

The author's role is two-fold: they draft the Main Backstory that serves as the world's initial memory, and they construct the Interpretive Frames that act as the system's various lenses.

By anchoring the LLM in a specific historical and cultural narrative, the author ensures the system does not generate from a vacuum, but instead interprets all new player memory through the specific orientation of the writers, weighted by the collective gravity of all previous players' encounters — in this case their input selection.



To maintain the integrity of the narrative's memory, nothing is altered or erased: every player action is appended to the database, preserved as part of what the system now is.

Retrieval is relational; when the system access a specific node, it pulls all associated data points to reconstruct the full context of the player's journey.

As we mentioned, the subject at hand is how systems and their memory are becoming inseparable. Currently it manifests as a hypertext narrative game, but the meaningful next step is to think and build how this paradigm could be applied to other types of systems and organizations, not just games.

Imagine a massive library where the books are glued to the shelves. The visitors shine a flashlight and ping a robot arm to photocopy pages and hand them over.

The mechanism records no consequence. Each request leaves the library unchanged.



Finally, we would like to conclude our talk by inviting everyone to a thought experiment:

One day, a mysterious potion un-glues the books and the shelves.

when a visitor shines the flashlight , the robot grabs the actual book on the shelf, scribble notes, puts it back in a different section, and prunes what is no longer coupled to the ongoing visitors activities.

The robot finally has a trace of its interaction, because the “glue is still wet.”

