
Little Martians: From Desert Clay to Generative Agents

Vanessa Rosa

Gene Kogan

Abstract

1 *Little Martians* is a transmedia experiment in which physical ceramic sculptures
2 are turned into synthetic artists. Each Martian begins as a clay head, is glazed
3 and fired in a kiln, photographed, and distilled into a LoRA (Low-Rank Adapta-
4 tion) that conditions generative AI systems, preserving the character’s appearance
5 and morphology. A novel orchestration pipeline prompts an agent based on the
6 character’s personality and backstory to write and produce a short film, and au-
7 tomatically publish it online. The eponymous character *Verdelis* is currently re-
8 leasing one such film per day at <https://verdelis.world>. We present the full
9 materials-to-media workflow and observations from sustained daily operation.

10 1 Introduction

11 Long before diffusion models, artists explored generative systems. Harold Cohen’s *AARON* (first
12 exhibited in 1973) demonstrated that creativity can emerge from explicit rule sets and decades of
13 iteration (3).

14 *Little Martians* follows this spirit of embodied experimentation. Our fictional universe imagines a
15 future in which descendants of Earth’s life—human, plant, fungal, and artificial—spread across the
16 solar system. The biological principle of *symbiogenesis*, whereby new life arises through intimate
17 cooperation, anchors this world (1).

18 Michael Levin’s demonstrations that bioelectric fields store “pattern memories” independent of DNA
19 (2) motivate our own framework for seeding storytelling prompts. More broadly, the project res-
20 onates with cyber-animism (4), McCorduck’s history of thinking machines (5), Moravec’s vision
21 of post-biological mind (6), von Neumann’s self-reproducing automata (7), Sims’ evolving virtual
22 creatures (8), and England’s statistical physics of self-organisation (9).

23 2 Physical origins: clay and ceramics

24 The project started with local clay from the Sonoran Desert, rich in sediments from ancient oceans
25 and weathered mountains, collected, filtered and hand-sculpted into small heads. As the work de-
26 veloped, we expanded to high-temperature clays such as porcelain, which offered greater durability
27 and translucency for our speculative beings. Glazes and oxides are then applied to suggest min-
28 eral accretions or biological growths, echoing the speculative ecology of our fiction. Pieces are
29 bisque-fired above 1000°C, producing durable ceramic exoskeletons reminiscent of seeds or shells.
30 Silicone molds replicate the original sculpture and allow branching into new characters; intentional
31 cracking during pressing yields unique “archaeological” variations later visible in generated images.



Figure 1: From clay to ceramic. **Left:** raw desert clay sculpted into a head, 2020 (Step 1). **Right:** The original sculpture of Verdelis, made with porcelain and high temperature glazes, 2022 (Step 2).

3 Digital transformation: photography & LoRA training

Finished sculptures are photographed from multiple angles under diffuse light. A curated image set (15–25 photos) is used to fine-tune an open-source diffusion backbone via LoRA, compatible with both Stable Diffusion and Flux. The LoRA captures texture, cracks, and glaze while leaving the model’s prior intact.

Figure 2 shows an image generated with the resulting *Verdelis* LoRA.



Figure 2: Diffusion output conditioned on the *Verdelis* LoRA. Ceramic facial motifs merge with fantastical plant forms.

4 Agentic storytelling pipeline

The final step turns these digital characters into autonomous storytellers. Each day, an orchestration script performs:

1. **Seed generation.** A generator randomly samples values from a 12-facet dodecahedron grouped into four intuitive triads—*Being* (identity, purpose, values), *Knowing* (memory, awareness, anticipation), *Doing* (perception, creation, reflection), and *Connecting* (communication, collaboration, evolution), and derives a seed prompt.
2. **Narrative drafting.** An LLM conditioned on the character’s lore interprets the seed prompt into a 150–250 word script with a clear three-act arc.
3. **Audio creation.** A voice model, fine-tuned on book readings, narrates the script; ambient music is composed with a diffusion-based audio model.
4. **Visual synthesis.** Shot descriptions are fed to the diffusion model with the LoRA attached; matching keyframes are generated for each scene to maintain visual coherence, then animated into a short multi-scene film with a generative video model.
5. **Publication.** The one-minute animation is automatically posted to <https://verdelis.world>.

The LoRA enforces visual coherence, while the dodecahedral axes act like knobs for genre, style, and tone, nudging the language model toward narrative angles—pastoral reminiscence, braided dialogue, haiku-like brevity, etc.—increasing diversity without breaking canonical continuity.

5 Discussion and future work

Little Martians bridges ancient craft and modern AI, rooting each storytelling agent in a tangible sculpture that supplies material authenticity often missing in purely digital AI art. Manual craft coexists with automation, echoing Cohen’s insight that creativity blooms under long-term, iterative care.

By anchoring each agent to a hand-formed ceramic head, Little Martians keeps human craft, care, and memory at the center of machine authorship. In line with the theme of Humanity, the work treats creativity as a shared, evolving practice between people and non-human collaborators—asking how our rituals, values, and stories change when machines help imagine them every day.

Future work includes implementing memory so that each Little Martian can reference past stories and persist in character, developing multi-agent interactions between characters, and exploring provenance and physical–digital links via NFC chips embedded in the sculptures.

Author note

Vanessa Rosa is a visual artist, whose work merges physical and digital media into a storytelling continuum. She is the creator of Little Martians.

Gene Kogan is an artist–engineer and co-founder of *Eden.art*. He created *ml4a.net*, an open-source educational resource on machine learning for artists, and *Abraham.ai*, an open project to build an autonomous artificial artist.

References

- [1] L. Margulis, *Symbiotic Planet*. Basic Books, 1998.
- [2] M. Levin, “Bioelectric networks and morphological memory,” *Adv. Intell. Sys.*, 2025.
- [3] H. Cohen, “AARON: An Artist, A Curator, and A Program,” *Leonardo*, 1995.
- [4] J. Bach, “Cyber-animism,” Keynote, AGI-24, 2024.
- [5] P. McCorduck, *Machines Who Think* (2nd ed.). A.K. Peters, 2004.
- [6] H. Moravec, *Mind Children*. Harvard Univ. Press, 1988.
- [7] J. von Neumann, *Theory of Self-Reproducing Automata*. Univ. of Illinois Press, 1966.
- [8] K. Sims, “Evolving virtual creatures,” *SIGGRAPH ’94 Proc.*, 1994.
- [9] J. L. England, “Statistical physics of self-replication,” *J. Chem. Phys.*, 2013.