

Simulacra

TL;DR - Try playing a game at <https://simulacra.cc>

An AI-powered tabletop exercise for crisis decision-making

Most AI risk discussion lives in blog posts and policy papers. You read about coordination failures, competing incentives, and misaligned objectives. You nod along. Then you close the tab and nothing changes.

Simulacra tries to make it experiential instead. It's a single-player strategy game where you role-play as a stakeholder during an escalating crisis. An LLM acts as the game master, generates the narrative, controls five AI opponents, and decides what your choices actually do to the world. You don't just read about how competing incentives cause coordination failures. You feel the pull of your own hidden objective while the shared public metric is dropping, and you make the tradeoff yourself. That's a different kind of understanding.

The name comes from Baudrillard. Simulacra are copies without originals, simulations that feel more real than reality. That's the conceit: you're playing through a synthetic crisis generated by an AI, making decisions alongside AI agents, and the whole thing still teaches you something about how real systems break. The simulation doesn't pretend to be reality. It just turns out to be useful anyway.

Can an LLM actually simulate a crisis?

On ForecastBench, GPT-4.5 hits a Brier score of 0.101 versus 0.081 for superforecasters. Not parity, but close, and the gap shrinks by about 0.016 points per year. For generating plausible "what happens next" scenarios in a game context, LLM world models are already solid. The bottleneck isn't prediction quality. It's making the experience engaging enough that people sit with the decisions instead of clicking through.

Where it's at

The stack is Next.js, React, TypeScript, Prisma, and PostgreSQL, with LLM calls routed through a LiteLLM proxy. The interesting engineering is in prompt design and action-tree generation.

Future work

- Multiplayer is mostly done. Fixing some bugs
- Grounded domain models (epidemiology, economics) would make specific scenarios more rigorous. But the core loop works
- counterfactual analysis after each round forces you to ask whether your clever move helped or just felt like it did.
- I am also working on adding resources and world models to improve the UX and design

The project is open source and looking for contributors.

Revision #4

Created 19 February 2026 10:59:20 by bhishma

Updated 9 March 2026 11:03:12 by bhishma