

Letter from the Editor-in-Chief

The central dilemma now confronting social computing is not simply technical. On one hand, human data is no longer sufficient. The corpora that sustained the past decade of AI advancement are depleted or contaminated, and the frontier of progress increasingly demands agents that learn by doing: by engaging, experimenting, and adapting in real time. This transition, described by Silver and Sutton as the onset of an *era of experience*, reorients machine learning away from static imitation and toward dynamic interaction with a changing world.

But on the other hand, the social web we rely on as a testbed for behavior, conversation, persuasion, and coordination is rapidly ceasing to be human. We are not sending agents into a reality composed purely of people. We are sending them into an ecosystem where, in short order, the majority of voices will be synthetic, the majority of interactions orchestrated by models, and the majority of “experience” recursively generated by other agents. We are entering an era where bots will train on the outputs of bots, converse with other bots, and optimize for outcomes shaped primarily by other bots’ reactions.

The promise of experiential learning was that it could escape the distortions and limitations of synthetic, crowd-sourced text, by grounding AI in consequence, not consensus. But what if the consequences themselves are generated by other models? A bot that learns to persuade by practicing on other bots is not becoming more socially intelligent; it is becoming more locally efficient within an artificial microeconomy. This raises the prospect of a new kind of collapse, not in the data, but in the learning process itself. A collapse not of signal, but of relevance.

What, then, can we do?

First, we must make provenance the central feature of every social dataset, algorithm, and evaluation benchmark. The distinction between human and synthetic actors cannot remain a soft annotation or post-hoc filter. It must shape the way we define cohesion, influence, virality, and trust. Metrics built for a human-only web will mislead us in a mixed-agent world, unless we explicitly model the ontological status of each node and edge. This includes not only identifying bots but understanding their architectures, their reward functions, and their deployment context. Not all synthetic actors are alike, and not all should be treated as distortive. But none should be treated as neutral.

Second, academia must invest in open, controlled, mixed-agent environments where human behavior and synthetic behavior can be studied in relation to one another. We lack the infrastructure to observe the long-term dynamics of humans and bots coexisting in a shared social substrate. Our benchmarks are either fully human (but outdated or polluted), or fully synthetic (and therefore unreal). We need instrumented social arenas, designed with consent, governance, and transparency, where agents learn from real-time interaction with humans, and where humans can push back, adapt, and shape the reward landscape. These are not just experiments; they are prototypes of the future.

Third, reward design must become a central concern of social computing. In the era of human data, we learned from what people had already said. In the era of experience, we must decide what outcomes matter. Do we reward an agent for maximizing engagement, or for preserving the diversity of viewpoints in a thread? Do we design for short-term persuasion, or for long-term trust? Grounded rewards are powerful, but they are also dangerous if they encode proxy metrics divorced from human flourishing. As AI begins to optimize social behavior directly, we must treat reward shaping as a sociotechnical act of governance.

Fourth, we must preserve, protect, and elevate human-authored data. As synthetic content floods the web, pre-2022 corpora and verifiable human discourse will become the rare earth minerals of the AI economy: scarce, valuable, and non-renewable. Academic institutions should take the lead in curating high-fidelity, consented, and demographically representative human data, not just as training material, but as a cultural archive. This is not nostalgia. It is epistemic infrastructure.

The papers in this issue offer glimpses of how these shifts are already underway. We see how cohesion

is being redefined, how community discovery is increasingly interactive and dynamic, how echo chambers and misinformation evolve in response to changing participation. We see experiments in influence-aware systems and visual abstractions that begin to accommodate complexity. Each paper speaks, in its own way, to the tension at the heart of our field: between structure and agency, between measurement and meaning, between simulation and experience.

But taken together, these contributions also remind us of our responsibility. Social computing was born as a way to understand human interaction at scale. That mandate does not disappear in the presence of synthetic actors. Instead, it becomes more urgent. If the future of intelligence is experiential, then the environments we build today will shape what intelligence becomes.

Let us ensure that what it learns is still meaningfully human.

Haixun Wang
EvenUp