Between the motion and the act

Chris von Csefalvay

Blog post published July 20, 2025 in *Chris von Csefalvay*
https://doi.org/10.59350/q3yj4-wah32

Abstract

There's a curious terminological schism that nobody talks about at what passes for dinner parties for the AI crowd.

Keywords

Epidemiology (Subfield), AI, Philosophy, LLMs, ABMs, Agentic AI

Recommended citation

von Csefalvay, C. (2025, July 20). Between the motion and the act. *Chris Von Csefalvay*. https://doi.org/10.59350/q3yj4-wah32

Copyright

Copyright © Chris von Csefalvay 2025. Distributed under the terms of the Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

There's a curious terminological schism that nobody talks about at what passes for dinner parties for the AI crowd. When Russell and Norvig crystallised the definition of an 'agent' in their 1995 book – "an entity that perceives its environment through sensors and acts upon it through actuators" –, they seemed to have a particular notion of agency that today's AI agents fit perfectly: they perceive prompts, process context and act by generating text, calling APIs or executing code. To those of us who cut our teeth on agent-based modelling, this is a little "yes, but". We have been cheerfully using the same word for something apparently different: simulated entities that exist only in computational Petri dishes, perceiving and acting upon nothing more substantial than bits and simulation spaces. One might reasonably ask whether we've been guilty of terminological theft, awkwardly appropriating a word that doesn't quite fit. But what if the divergence is an illusion? What if simulation and action aren't opposing categories but different points on the same spectrum of agency?

Everything old is new again

This is the way of all technological revolutions.[1] We become so intoxicated by the new that we forget the profound insights buried in what came before. It's rather like discovering molecular gastronomy and suddenly forgetting that your grandma's stock pot held secrets that no amount of liquid nitrogen could replicate. I've spent considerable time with agent-based models – enough to devote an entire chapter to them in *Computational Modeling of Infectious Disease* –, watching imaginary pathogens spread through synthetic populations.[2] These populations existed purely in silico, bumping into each other according to rules we painstakingly crafted to mirror the messiness of human behaviour. Beautiful, horrifying and above all, instructive.

[1] I will to my dying day resist calling it a *paradigm shift*. A paradigm shift is when your mental model changes. A revolution is when the world changes. A guillotine is what happens when the second of these occurs without you responding timely enough with the first.

[2] It's been a perverse pleasure of mine to do so over places I knew well. My book has a geospatial simulation of RESTV, which would have rather early reached my home, less than 10mi from the epicentre.

These agents weren't trying to book your flights or write your emails. They were thinking machines in the truest sense – not because they possessed consciousness, but because they allowed us to think through them. They acted as our cognitive tools, extending our ability to reason about complex systems by playing out scenarios we couldn't possibly compute in our heads (never mind simulate in a lab). To the rather less abstract-disposed AI community, agents meant more. They could be more than passive observers: they could be actors in the world, semi-embodied (for now), commanding not physical bodies but prehensile digital limbs of APIs, MCP calls and A2A connections.

This split felt natural, even inevitable. After all, why merely simulate when you could actuate? Why watch synthetic populations of stock brokers trade imaginary shares when you could build agents that trade real ones and get silly rich in an afternoon? The logic seemed unassailable. But we've

been victims of a false dichotomy – treating simulation and action as opposing categories rather than complementary modes of agency that desperately need reunification: for the sake of both domains.

Simulaction

What I'm proposing isn't just a nostalgic return to ABM. It's a reconciliation, a synthesis. Imagine agent ecosystems where contemplative swarms continuously simulate possible futures, while their active cousins execute in the present. The simulators become the dreamers, the actors become the hands.

This isn't as far-fetched as it might sound. In my own work in computational epidemiology, we've long used simulations to inform policy. But these were primarily disconnected processes – we'd run our models, generate our reports and hope someone would read them. What if, instead, the simulation agents could directly communicate with action agents?

Picture this: a swarm of ABM agents continuously simulating supply chain dynamics, playing out scenarios of disruption, adaptation and recovery. When they converge on a particularly robust finding – say, a vulnerability in a specific shipping route – they don't generate a `ggplot`. They A2A an action agent that can address it.

The possibility machine

The concept of world models in robotics offers us a profound insight: simulation isn't just about prediction, it's about exploration. When a robot in Isaac Sim attempts a thousand different ways to grasp an object, it's not trying to predict which one will work. It's building an experiential understanding of the entire possibility space. This is precisely what we need in agentic AI systems: not agents that try to predict the future (that way lie madness and margin calls) but agents that understand the shape of possible futures. They need to know not just what might happen, but how different actions change the topology of this possibility space.

In practice: a financial services firm might deploy simulation agents not to predict tomorrow's stock prices but to understand how different market conditions interact. These agents would run thousands of parallel explorations of the decision space. What happens if interest rates rise while supply chains remain constrained? What if consumer confidence drops but corporate earnings stay strong?

The active trading agents would then use the carefully mapped picture of the possiblity space and know which actions open up more options, which close them down, which create resilience and which create fragility. In turn, the real-world outcomes as perceived and reported by the active trading agents would ground the simulation's futures (no point in simulating what doesn't work). Did that supply chain intervention help? Was their call to follow earnings reports over consumer sentiment the right one? The simulation agents need to know how to update their models – how to dream better dreams next time.

https://doi.org/10.59350/q3yj4-wah32

The practical implications of this 'possiblity machine' are of course technologically challenging. The nitpickers (and I say that with love – pedantry is just rigour in a bow tie) will immediately spot them: how do we ensure consistency between the simulated worlds and the actual world? How do we implement that feedback? How do we create guardrails without excess transparency? How are we going to pay for all this?

I remain sanguine about the fusion of simulation and action for two reasons, despite acknowledging the validity of each and every one of these concerns. For one, these are technical concerns, and those tend to get solved as time converges to whatever it converges towards. On the other hand, they're also indicative of bigger questions, and that always hints at there being not a wall but a wardrobe-gate to some hitherto unexplored aspects of reality. These concerns ultimately are fundamental questions about the nature of agency, simulation and reality. To me at least, they point the way towards the simulation/action paradigm's version of the question of the agentic agora: how we move from message passing to shared state spaces (some of which might even be real).

The dream of the swarm

At no point am I talking about a notion of what the 'right' job is for agents, simulation or action. Rather, I'd like to see both coexist in productive tension, each making the other more effective. The simulation swarms dream of a diverging universe of futures that action agents sample, enact and critique in view of their perception of reality. Like the neighbour's annoying lawn mower on a Sunday afternoon, those observations filter back into the dreams. Something mostly akin to insight, maybe even wisdom, emerges – of the kind that neither breed of agents with a limited purview could on its own attain.

The agentic revolution gave us agents that could do. ABMs gave us agents that can dream – and more importantly, agents that can learn from those dreams. And in the space between dreaming and doing, we may find a new paradigm that unifies the entire OODA loop[3] into a coherent framework.

[3] Yes, I *know* it's not a 'loop', strictly speaking. .

Or at least better supply chains. I'd settle for that.

https://doi.org/10.59350/q3yj4-wah32